

QUADRATIC DIFFERENTIALS AS STABILITY CONDITIONS

TOM BRIDGELAND AND IVAN SMITH

ABSTRACT. We prove that moduli spaces of meromorphic quadratic differentials with simple zeroes on compact Riemann surfaces can be identified with spaces of stability conditions on a class of CY_3 triangulated categories defined using quivers with potential associated to triangulated surfaces. We relate the finite-length trajectories of such quadratic differentials to the stable objects of the corresponding stability condition.

CONTENTS

1. Introduction	1
2. Quadratic differentials	13
3. Trajectories and geodesics	22
4. Period co-ordinates	32
5. Stratification by number of separating trajectories	42
6. Colliding zeroes and poles: the spaces $\text{Quad}(\mathbb{S}, \mathbb{M})$	56
7. Quivers and stability conditions	64
8. Surfaces and triangulations	75
9. The category associated to a surface	84
10. From differentials to stability conditions	93
11. Proofs of the main results	100
12. Examples	112
References	121

1. INTRODUCTION

In this paper we prove that spaces of stability conditions on a certain class of triangulated categories can be identified with moduli spaces of meromorphic quadratic differentials. The relevant categories are Calabi-Yau of dimension three (CY_3), and are described using quivers with potential associated to triangulated surfaces. The observation that spaces of abelian and quadratic differentials have similar properties to spaces of stability conditions was first made by Kontsevich and Seidel several years ago. On the one hand, our results provide some of the first descriptions of spaces of stability conditions on CY_3 categories, which is the case of most interest in physics. On the other, they give a precise link between the trajectory structure of flat surfaces and the theory of wall-crossing and Donaldson-Thomas invariants.

T.B. is supported by All Souls College, Oxford.

I.S. is partially supported by a grant from the European Research Council.

Our results can also be viewed as a first step towards a mathematical understanding of the work of physicists Gaiotto, Moore and Neitzke [13, 14]. Their paper [13] describes a remarkable interpretation of the Kontsevich-Soibelman wall-crossing formula for Donaldson-Thomas invariants in terms of hyperkähler geometry. In the sequel [14] an extended example is described, relating to parabolic Higgs bundles of rank two. The mathematical objects studied in the present paper are very closely related to their physical counterparts in [14], and some of our basic constructions are taken directly from that paper. We hope to return to the relations with Hitchin systems and cluster varieties in a future publication. In another direction, the CY_3 categories appearing in this paper also arise as Fukaya categories of certain quasi-projective Calabi-Yau threefolds. That relation is the subject of a sequel [32] currently in preparation.

In this introductory section we shall first recall some basic facts about quadratic differentials on Riemann surfaces. We then describe the simplest examples of the categories we shall be studying, before giving a summary of our main result in that case, together with a very brief sketch of how it is proved. We then state the other version of our result involving quadratic differentials with higher-order poles. We conclude by discussing the relationship between the finite-length trajectories of a quadratic differential and the stable objects of the corresponding stability condition.

As a matter of notation, the triangulated categories we consider here are most naturally labelled by combinatorial data consisting of a smooth surface \mathbb{S} equipped with a collection of marked points $\mathbb{M} \subset \mathbb{S}$, all considered up to diffeomorphism. Initially \mathbb{S} will be closed, but in the second form of our result \mathbb{S} can have non-empty boundary. The quadratic differentials we consider live on Riemann surfaces S whose underlying smooth surface is obtained from \mathbb{S} by collapsing each boundary component to a point. To avoid confusion, we shall try to preserve the notational distinction whereby \mathbb{S} refers to a smooth surface, possibly with boundary, whereas S is always a Riemann surface, usually compact. All these surfaces will be assumed to be connected.

We fix an algebraically closed field k throughout.

1.1. Quadratic differentials. A meromorphic quadratic differential ϕ on a Riemann surface S is a meromorphic section of the holomorphic line bundle $\omega_S^{\otimes 2}$. We emphasize that all the differentials considered in this paper will be assumed to have simple zeroes. Two quadratic differentials ϕ_1, ϕ_2 on Riemann surfaces S_1, S_2 are considered to be equivalent if there is a holomorphic isomorphism $f: S_1 \rightarrow S_2$ such that $f^*(\phi_2) = \phi_1$.

Let \mathbb{S} be a compact, closed, oriented surface, with a non-empty set of marked points $\mathbb{M} \subset \mathbb{S}$. We assume that if $g(\mathbb{S}) = 0$ then $|\mathbb{M}| \geq 3$. Up to diffeomorphism the pair (\mathbb{S}, \mathbb{M}) is determined by the genus $g = g(\mathbb{S})$ and the number $d = |\mathbb{M}| > 0$ of marked points. We use this combinatorial data to specify a union of strata in the space of meromorphic quadratic differentials; this will be less trivial later when we allow \mathbb{S} to have boundary.

By a quadratic differential on (\mathbb{S}, \mathbb{M}) we shall mean a pair (S, ϕ) , where S is a compact and connected Riemann surface of genus $g = g(\mathbb{S})$, and ϕ is a meromorphic quadratic differential with simple zeroes and exactly $d = |\mathbb{M}|$ poles, each one of order ≤ 2 . Note that every equivalence class of such differentials contains pairs (S, ϕ) such that \mathbb{S} is the underlying smooth surface of S , and ϕ has poles precisely at the points of \mathbb{M} .

A quadratic differential (S, ϕ) of this form determines a double cover $\pi: \hat{S} \rightarrow S$, called the spectral cover, branched precisely at the zeroes and the simple poles of ϕ . This cover has the property that

$$\pi^*(\phi) = \psi \otimes \psi$$

for some globally-defined meromorphic 1-form ψ . We write $\hat{S}^\circ \subset \hat{S}$ for the complement of the poles of ψ . The hat-homology group of the differential (S, ϕ) is defined to be

$$\hat{H}(\phi) = H_1(\hat{S}^\circ; \mathbb{Z})^-$$

where the superscript indicates the anti-invariant part for the action of the covering involution. The 1-form ψ is holomorphic on \hat{S}° and anti-invariant, and hence defines a de Rham cohomology class, called the period of ϕ , which we choose to view as a group homomorphism

$$Z_\phi: \hat{H}(\phi) \rightarrow \mathbb{C}, \quad \gamma \mapsto \int_\gamma \psi.$$

There is a complex orbifold $\text{Quad}(\mathbb{S}, \mathbb{M})$ of dimension

$$n = 6g - 6 + 3d$$

parameterizing equivalence-classes of quadratic differentials on (\mathbb{S}, \mathbb{M}) . We call a quadratic differential complete if it has no simple poles; such differentials form a dense open subset $\text{Quad}(\mathbb{S}, \mathbb{M})_0 \subset \text{Quad}(\mathbb{S}, \mathbb{M})$.

The homology groups $\hat{H}(\phi)$ form a local system over the orbifold $\text{Quad}(\mathbb{S}, \mathbb{M})_0$. A slightly subtle point is that this local system does not extend over $\text{Quad}(\mathbb{S}, \mathbb{M})$, but rather has monodromy of order 2 around each component of the divisor parameterizing differentials with a simple pole. It therefore defines a local system on an orbifold $\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$ which has larger automorphism groups along this divisor. There is a natural map

$$\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) \rightarrow \text{Quad}(\mathbb{S}, \mathbb{M}),$$

which is an isomorphism over the open subset $\text{Quad}(\mathbb{S}, \mathbb{M})_0$, and which induces an isomorphism on coarse moduli spaces. Fixing a free abelian group Γ of rank n , we can also consider an unramified cover

$$\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M}) \longrightarrow \text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$$

of framed quadratic differentials, consisting of equivalence classes of quadratic differentials as above, equipped with a local trivialization $\Gamma \cong \hat{H}(\phi)$ of the hat-homology local system.

In Section 4 we shall prove the following result, which is a variation on the usual existence of period co-ordinates in spaces of quadratic differentials. For this we need to assume that (\mathbb{S}, \mathbb{M}) is not a torus with a single marked point.

Theorem 1.1. *The space of framed differentials $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$ is a complex manifold, and there is a local homeomorphism*

$$(1.1) \quad \pi: \text{Quad}^\Gamma(\mathbb{S}, \mathbb{M}) \longrightarrow \text{Hom}_{\mathbb{Z}}(\Gamma, \mathbb{C}),$$

obtained by composing the framing and the period.

In the excluded case the space $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$ is not a manifold because it has generic automorphism group \mathbb{Z}_2 .

1.2. Triangulations and quivers. Suppose again that \mathbb{S} is a compact, closed, oriented surface with a non-empty set of marked points $\mathbb{M} \subset \mathbb{S}$. For the purposes of the following discussion we will assume that if $g(\mathbb{S}) = 0$ then $|\mathbb{M}| \geq 5$.

By a non-degenerate ideal triangulation of (\mathbb{S}, \mathbb{M}) we mean a triangulation of \mathbb{S} whose vertex set is precisely \mathbb{M} and in which every vertex has valency at least 3. To each such triangulation T there is an associated quiver $Q(T)$ whose vertices are the midpoints of the edges of T , and whose arrows are obtained by inscribing a small clockwise 3-cycle inside each face of T , as illustrated in Figure 1.




FIGURE 1. Quiver associated to a triangulation.

There are two obvious systems of cycles in $Q(T)$, namely a clockwise 3-cycle $T(f)$ in each face f , and an anticlockwise cycle $C(p)$ of length at least 3 encircling each point $p \in \mathbb{M}$. We define a potential $W(T)$ on $Q(T)$ by taking the sum

$$W(T) = \sum_f T(f) - \sum_p C(p).$$

Consider the derived category of the complete Ginzburg algebra [15, 20] of the quiver with potential $(Q(T), W(T))$ over k , and let $\mathcal{D}(T)$ be the full subcategory consisting of modules with finite-dimensional cohomology. It is a CY_3 triangulated category of finite type over k , and comes equipped with a canonical t-structure, whose heart $\mathcal{A}(T) \subset \mathcal{D}(T)$ is equivalent to the category of finite-dimensional modules for the completed Jacobi algebra of $(Q(T), W(T))$.

Suppose that two non-degenerate ideal triangulations T_i are related by a flip, in which the diagonal of a quadrilateral is replaced by its opposite diagonal, as in Figure 2. The point of the above definition is that the resulting quivers with potential $(Q(T_i), W(T_i))$ are related by a mutation at the vertex corresponding to the edge being flipped; see Figure 2. It follows from general results of Keller and Yang [20] that there is a distinguished pair of k -linear triangulated equivalences $\Phi_{\pm}: \mathcal{D}(T_1) \cong \mathcal{D}(T_2)$.




FIGURE 2. Effect of a flip.

Labardini-Fragoso [24] extended the correspondence between ideal triangulations and quivers with potential so as to encompass a larger class of triangulations containing vertices of valency ≤ 2 . He then proved the much more difficult result that flips induce mutations in this more general context. Since any two of these more general ideal triangulations are related by a finite chain of flips, it follows that up to k -linear triangulated equivalence, the category $\mathcal{D}(T)$ is independent of the chosen triangulation. We loosely use the notation $\mathcal{D}(\mathbb{S}, \mathbb{M})$ to denote any triangulated category $\mathcal{D}(T)$ defined by an ideal triangulation T of the marked surface (\mathbb{S}, \mathbb{M}) .

1.3. Stability conditions. A stability condition on a triangulated category \mathcal{D} is a pair $\sigma = (Z, \mathcal{P})$ consisting of a group homomorphism $Z: K(\mathcal{D}) \rightarrow \mathbb{C}$ called the central charge, and an \mathbf{R} -graded collection of objects

$$\mathcal{P} = \bigcup_{\phi \in \mathbf{R}} \mathcal{P}(\phi) \subset \mathcal{D}$$

known as the semistable objects, which together satisfy some axioms (see Section 7.5).

For simplicity, let us assume that the Grothendieck group $K(\mathcal{D})$ is free of some finite rank n . There is then a complex manifold $\text{Stab}(\mathcal{D})$ of dimension n whose points are stability conditions on \mathcal{D} satisfying a further condition known as the support property. The map

$$(1.2) \quad \pi: \text{Stab}(\mathcal{D}) \longrightarrow \text{Hom}_{\mathbb{Z}}(K(\mathcal{D}), \mathbb{C})$$

taking a stability condition to its central charge is a local homeomorphism. The manifold $\text{Stab}(\mathcal{D})$ carries a natural action of the group $\text{Aut}(\mathcal{D})$ of triangulated autoequivalences of \mathcal{D} .

Now suppose that (\mathbb{S}, \mathbb{M}) is a compact, closed, oriented surface with marked points, and let \mathcal{D} be the CY_3 triangulated category $\mathcal{D}(\mathbb{S}, \mathbb{M})$ defined in the last subsection. There is a distinguished connected component

$$\text{Stab}_\Delta(\mathcal{D}) \subset \text{Stab}(\mathcal{D}),$$

containing stability conditions whose heart is one of the standard hearts $\mathcal{A}(T) \subset \mathcal{D}(T)$ discussed above. We write

$$\text{Aut}_\Delta(\mathcal{D}) \subset \text{Aut}(\mathcal{D})$$

for the subgroup of autoequivalences of \mathcal{D} which preserve this component. We also define $\mathcal{A}ut_\Delta(\mathcal{D})$ be the quotient of $\text{Aut}_\Delta(\mathcal{D})$ by the subgroup of autoequivalences which act trivially on $\text{Stab}_\Delta(\mathcal{D})$.

The first form of our main result is

Theorem 1.2. *Let (\mathbb{S}, \mathbb{M}) be a compact, closed, oriented surface with marked points. Assume that one of the following two conditions holds*

- (a) $g(\mathbb{S}) = 0$ and $|\mathbb{M}| > 5$;
- (b) $g(\mathbb{S}) > 0$ and $|\mathbb{M}| > 1$.

Then there is an isomorphism of complex orbifolds

$$\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) \cong \text{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta(\mathcal{D}).$$

The assumption on the number of punctures in the $g(\mathbb{S}) = 0$ case of Theorem 1.2 comes from a similar restriction in a crucial result of Labardini-Fragoso [26]. We conjecture that the conclusion of the Theorem holds with the weaker assumptions that $|\mathbb{M}| > 1$ and that if $g(\mathbb{S}) = 0$ then $|\mathbb{M}| > 3$. The case of a once-punctured surface is special in many respects, and we leave it for future research; see Section 11.6 for more comments on this. The case of a three-punctured sphere is also special, and is treated in Section 12.4.

1.4. Horizontal strip decomposition. The main ingredient in the proof of Theorem 1.2 is the statement that a generic point of the space $\text{Quad}(\mathbb{S}, \mathbb{M})$ determines an ideal triangulation of the surface (\mathbb{S}, \mathbb{M}) , well-defined up to the action of the mapping class group. We learnt this idea from Gaiotto, Moore and Neitzke's work [14, Section 6], although in retrospect, it is an immediate consequence of well-known results in the theory of quadratic differentials.

Away from its zeroes and poles, a quadratic differential ϕ on a Riemann surface S induces a flat metric, together with a foliation known as the horizontal foliation. One

way to see this is to write $\phi = dz^{\otimes 2}$ for some local co-ordinate z , well-defined up to $z \mapsto \pm z + \text{constant}$. The metric is then given by pulling back the Euclidean metric on \mathbb{C} using z , and the horizontal foliation is given by the lines $\text{Im}(z) = \text{constant}$.

The integral curves of the horizontal foliation are called trajectories. The trajectory structure near a simple zero and a generic double pole are illustrated in Figure 3. Note




FIGURE 3. Local trajectory structure at a simple zero and a generic double pole.

that generic double poles behave like black holes: any trajectory passing beyond a certain event horizon eventually falls into the pole. Thus for a generic differential one expects all trajectories to tend towards a double pole in at least one direction.

In the flat metric on S induced by ϕ , any pole of order ≥ 2 lies at infinity. Therefore, assuming that S is compact, any finite-length trajectory γ is either a simple closed curve containing no critical points of ϕ , or is a simple arc which tends to a finite critical points of ϕ (zeroes or simple poles) at either end. In the first case γ is called a closed trajectory, and moves in an annulus of such trajectories known as a ring domain. In the second case we call γ a saddle trajectory. Note that the endpoints of a saddle trajectory γ could well coincide; when this happens we call γ a closed saddle trajectory.

The boundary of a ring domain has two components, and each usually consists of unions of saddle trajectories. There is one other possibility however: a ring domain may consist of closed curves encircling a double pole p with real residue; the point p is then one of the boundary components. We call such ring domains degenerate.




FIGURE 4. A saddle trajectory, a ring domain and a degenerate ring domain.

There is a dense open subset $B_0 \subset \text{Quad}(\mathbb{S}, \mathbb{M})$ consisting of differentials (S, ϕ) with no simple poles and no finite-length trajectories; we call such differentials saddle-free. For saddle-free differentials, each of the three horizontal trajectories leaving a given

zero eventually tend towards a double pole. These separating trajectories divide the surface S into a union of cells, known as horizontal strips (see Figure 5). Taking a




FIGURE 5. The separating (solid) and generic trajectories (dotted) for a saddle-free differential; the black dots represent double poles.

single generic trajectory from each horizontal strip gives a triangulation of the surface S , whose vertices lie at the poles of ϕ , and this then induces an ideal triangulation T of the surface (\mathbb{S}, \mathbb{M}) , well-defined up to the action of the mapping class group. This is what is referred to as the WKB triangulation in [14].

The dual graph to the collection of separating trajectories is precisely the quiver $Q(T)$ considered before. In particular, the vertices of $Q(T)$ naturally correspond to the horizontal strips of ϕ . In each horizontal strip h_i there is a unique homotopy class of arcs ℓ_i joining the two zeroes of ϕ lying on its boundary. Lifting ℓ_i to the spectral cover gives a class $\alpha_i \in \hat{H}(\phi)$, and taken together, these classes form a basis. There is thus a natural isomorphism

$$\nu: K(\mathcal{D}(T)) \rightarrow \hat{H}(\phi),$$

which sends the class of the simple module S_i at a vertex of $Q(T)$, to the class α_i defined by the corresponding horizontal strip h_i .

Using the isomorphism ν , the period of ϕ can be interpreted as a group homomorphism $Z_\phi: K(\mathcal{D}(T)) \rightarrow \mathbb{C}$. More concretely, this is given by

$$Z_\phi(S_i) = 2 \int_{\ell_i} \sqrt{\phi} \in \mathbb{C},$$

where the sign of $\sqrt{\phi}$ is chosen so that $\text{Im } Z_\phi(S_i) > 0$. We thus have a triangulated category $\mathcal{D}(T)$, with its canonical heart $\mathcal{A}(T)$, and a compatible central charge Z_ϕ . This is precisely the data needed to define a stability condition on $\mathcal{D}(T)$.

We refer to the connected components of the open subset B_0 as chambers; the horizontal strip decomposition and the triangulation T are constant in each chamber, although

the period Z_ϕ varies. As one moves from one chamber to a neighbouring one, the triangulation T can undergo a flip. Gluing the stability conditions obtained from all these chambers using the Keller-Yang equivalences Φ_\pm referred to above eventually leads to a proof of Theorem 1.2.

1.5. Higher-order poles. We can extend Theorem 1.2 to cover quadratic differentials with poles of order > 2 . Such differentials correspond to stability conditions on categories defined by triangulations of surfaces with boundary. For this reason it will be convenient to also index the relevant moduli spaces of differentials by such surfaces, as we now explain.

A marked, bordered surface (\mathbb{S}, \mathbb{M}) is a pair consisting of a compact, oriented, smooth surface \mathbb{S} , possibly with boundary, together with a collection of marked points $\mathbb{M} \subset \mathbb{S}$, such that every boundary component of \mathbb{S} contains at least one point of \mathbb{M} . The marked points $\mathbb{P} \subset \mathbb{M}$ lying in the interior of \mathbb{S} are called punctures. We shall always assume that (\mathbb{S}, \mathbb{M}) is not one of the following:

- (i) a sphere with ≤ 2 punctures;
- (ii) an unpunctured disc with ≤ 2 marked points on its boundary.

These excluded surfaces have no ideal triangulations, and hence our theory is vacuous.

The trajectory structure of a quadratic differential ϕ near a higher-order pole is illustrated in Figure 6; just as with double poles there is an event horizon beyond which all trajectories tend to the pole, but at a pole of order $k + 2$ there are, in addition, k distinguished tangent vectors along which all trajectories enter.




FIGURE 6. Local trajectory structure at a pole of order 5.

A meromorphic quadratic differential ϕ on a compact Riemann surface S determines a marked, bordered surface (\mathbb{S}, \mathbb{M}) by the following construction. To define the surface \mathbb{S} we take the underlying smooth surface of S and perform an oriented real blow-up at each pole of ϕ of order ≥ 3 . The marked points \mathbb{M} are then the poles of ϕ of order ≤ 2 , considered as points of the interior of \mathbb{S} , together with the points on the boundary of \mathbb{S} corresponding to the distinguished tangent directions.

Let us now fix a marked, bordered surface (\mathbb{S}, \mathbb{M}) . Let $\text{Quad}(\mathbb{S}, \mathbb{M})$ denote the space of equivalence classes of pairs (S, ϕ) , consisting of a compact Riemann surface S , together with a meromorphic quadratic differential ϕ with simple zeroes, whose associated marked bordered surface is diffeomorphic to (\mathbb{S}, \mathbb{M}) .

More concretely, the pair (\mathbb{S}, \mathbb{M}) is determined up to diffeomorphism by the genus $g = g(\mathbb{S})$, the number of punctures $p = |\mathbb{P}|$, and a collection of integers $k_i \geq 1$ encoding the number of marked points on each boundary component of \mathbb{S} . The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ then consists of equivalence classes of pairs (S, ϕ) consisting of a meromorphic quadratic differential ϕ on a compact Riemann surface S of genus g , having p poles of order ≤ 2 , a collection of higher-order poles with multiplicities $k_i + 2$, and simple zeroes.

The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ is a complex orbifold of dimension

$$n = 6g - 6 + 3p + \sum_i (k_i + 3).$$

We can define the spectral cover $\pi: \hat{S} \rightarrow S$, the hat-homology group $\hat{H}(\phi)$, and the spaces $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$ and $\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$ exactly as before. We can also prove the analogue of Theorem 1.1 in this more general setting.

The theory of ideal triangulations of marked bordered surfaces has been developed for example in [10]. The results of Labardini-Fragoso [24] apply equally well in this more general situation, so exactly as before, there is a CY_3 triangulated category $\mathcal{D} = \mathcal{D}(\mathbb{S}, \mathbb{M})$, well-defined up to k -linear equivalence, and a distinguished connected component $\text{Stab}_\Delta(\mathcal{D})$.

The second form of our main result is

Theorem 1.3. *Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface with non-empty boundary. Then there is an isomorphism of complex orbifolds*

$$\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) \cong \text{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta(\mathcal{D}).$$

There are six degenerate cases which have been suppressed in the statement of Theorem 1.3. Firstly, if (\mathbb{S}, \mathbb{M}) is one of the following three surfaces

- (a) a once-punctured disc with 2 or 4 marked points on the boundary;
- (b) a twice-punctured disc with 2 marked points on the boundary;

then Theorem 1.3 continues to hold, but only if we replace $\mathcal{A}ut_\Delta(\mathcal{D})$ by a certain index 2 subgroup $\mathcal{A}ut_\Delta^{\text{allow}}(\mathcal{D})$. The basic reason for this is that a triangulation T of such a surface is not determined up to the action of the mapping class group by the associated quiver $Q(T)$. Secondly, if (\mathbb{S}, \mathbb{M}) is one of the following three surfaces

- (c) an unpunctured disc with 3 or 4 marked points on the boundary;

(d) an annulus with one marked point on each boundary component;

then the space $\text{Quad}(\mathbb{S}, \mathbb{M})$ has a generic automorphism group which must first be killed to make Theorem 1.3 hold. See Section 11.6 for more details on these exceptional cases.

Particular choices of the data (\mathbb{S}, \mathbb{M}) lead to quivers of interest in representation theory. See Section 12 for some examples of this. In particular, we can recover in this way some recent results of T. Sutherland [34, 35], who used different methods to compute the spaces of numerical stability conditions on the categories $\mathcal{D}(\mathbb{S}, \mathbb{M})$ in all cases in which these spaces are two-dimensional.

1.6. Saddle trajectories and stable objects. In the course of proving the Theorems stated above, we will in fact prove a stronger result, which gives a direct correspondence between the finite-length trajectories of a quadratic differential and the stable objects of the corresponding stability condition.

To describe this correspondence in more detail, fix a marked bordered surface (\mathbb{S}, \mathbb{M}) satisfying the assumptions of one of our main theorems, and let $\mathcal{D} = \mathcal{D}(\mathbb{S}, \mathbb{M})$ be the corresponding triangulated category. Let ϕ be a meromorphic differential on a compact Riemann surface S defining a point $\phi \in \text{Quad}(\mathbb{S}, \mathbb{M})$, and let $\sigma \in \text{Stab}(\mathcal{D})$ be the corresponding stability condition, well-defined up to the action of the group $\text{Aut}_\Delta(\mathcal{D})$. We shall say that the differential ϕ is generic if for any two hat-homology classes $\gamma_i \in \hat{H}(\phi)$

$$\mathbf{R} \cdot Z_\phi(\gamma_1) = \mathbf{R} \cdot Z_\phi(\gamma_2) \implies \mathbb{Z} \cdot \gamma_1 = \mathbb{Z} \cdot \gamma_2.$$

Generic differentials form a dense subset of $\text{Quad}(\mathbb{S}, \mathbb{M})$, and for simplicity we shall restrict our attention to these.

To state the result, let us denote by $\mathcal{M}_\sigma(0)$ the moduli space of objects in \mathcal{D} that are stable in the stability condition σ and of phase 0. This space can be identified with a moduli space of stable representations of a finite-dimensional algebra, and hence by work of King [21], is represented by a quasi-projective scheme over k .

Theorem 1.4. *Assume that ϕ is generic. Then $\mathcal{M}_\sigma(0)$ is smooth, and each of its connected components is either a point, or is isomorphic to the projective line \mathbb{P}^1 . Moreover, there are bijections*

$$\begin{aligned} \{0\text{-dimensional components of } \mathcal{M}_\sigma(0)\} &\longleftrightarrow \{\text{non-closed saddle trajectories of } \phi\}; \\ \{1\text{-dimensional components of } \mathcal{M}_\sigma(0)\} &\longleftrightarrow \{\text{non-degenerate ring domains of } \phi\}. \end{aligned}$$

Note that with our conventions, all trajectories are assumed to be horizontal, and correspond to stable objects of phase 0. In particular, a stability condition σ has a stable object of phase 0 precisely if the corresponding differential ϕ has a finite-length trajectory. Stable objects of more general phases θ correspond in exactly the same way

to finite-length straight arcs which meet the horizontal foliation at a constant angle $\pi\theta$. This more general statement follows immediately from Theorem 1.4, because the isomorphisms of our main theorems are compatible with the natural \mathbb{C}^* -actions on both sides.

Standard results in Donaldson-Thomas theory imply that the two types of moduli spaces appearing in Theorem 11.6 contribute $+1$ and -2 respectively to the BPS invariants, although we do not include the proof of this here. These exactly match the contributions to the BPS invariants described in [14, Section 7.6]. In physics terminology, non-closed saddle trajectories correspond to BPS hypermultiplets, and non-degenerate ring domains to BPS vectormultiplets.

It is a standard open question in the theory of flat surfaces to characterise or constrain the hat-homology classes which contain saddle connections. Theorem 1.4 relates this to the similar problem of identifying the classes in the Grothendieck group which support stable objects. Here one has the powerful technology of Donaldson-Thomas invariants and the Kontsevich-Soibelman wall-crossing formula [23], which in principle allows one to determine how the spectrum of stable objects changes as the stability condition varies. It would be interesting to see whether these techniques can be usefully applied to the theory of flat surfaces.

1.7. Structure of the paper. The paper splits naturally into three parts.

The first part, consisting of Sections 2–6, is concerned with spaces of meromorphic quadratic differentials. Section 2 reviews basic notions concerning quadratic differentials, and introduces orbifolds $\text{Quad}(g, m)$ parameterizing differentials with simple zeroes and fixed pole orders. Section 3 consists of well-known material on the trajectory structure of quadratic differentials. Section 4 is devoted to proving that the period map on $\text{Quad}(g, m)$ is a local isomorphism. Section 5 studies the stratification of the space $\text{Quad}(g, m)$ by the number of separating trajectories. Finally, Section 6 introduces the spaces $\text{Quad}(\mathbb{S}, \mathbb{M})$ appearing above, in which zeroes of the differentials are allowed to collide with the double poles.

The second part, comprising Sections 7–9, is concerned with CY_3 triangulated categories, and more particularly, the categories $\mathcal{D}(\mathbb{S}, \mathbb{M})$ described above. Section 7 consists of general material on quivers with potential, t-structures, tilting and stability conditions. Section 8 introduces the basic combinatorial properties of ideal and tagged triangulations. Section 9 contains a more detailed study of the categories $\mathcal{D}(\mathbb{S}, \mathbb{M})$, including their autoequivalence groups, and gives a precise correspondence relating t-structures on $\mathcal{D}(\mathbb{S}, \mathbb{M})$ to tagged triangulations of the surface (\mathbb{S}, \mathbb{M}) .

The geometry and algebra come together in the last part, which comprises Sections 10–12. Section 10 describes the WKB triangulation associated to a saddle-free differential, and the way it changes as one passes between neighbouring chambers. Section 11

contains the proofs of our main results identifying spaces of stability conditions with spaces of quadratic differentials. We finish in Section 12 with some illustrative examples.

The reader is advised to start with §§2–3, the first half of §6, and §§7–9, since these contain the essential definitions and are the least technical. It may also help to look at some of the examples in §12.

Acknowledgements. Thanks most of all to Daniel Labardini-Fragoso, Andy Neitzke and Tom Sutherland, all of whom have been enormously helpful. Thanks too to Sergey Fomin, Bernhard Keller, Alastair King, Howard Masur, Michael Shapiro and Anton Zorich for helpful conversations and correspondence. This paper owes a significant debt to the work of Davide Gaiotto, Greg Moore and Andy Neitzke [14].

2. QUADRATIC DIFFERENTIALS

We begin by summarizing some of the basic properties of meromorphic quadratic differentials on Riemann surfaces. This material is mostly well-known, although we were unable to find any references dealing with the moduli spaces of differentials with higher-order poles that we shall be using. Our standard reference for quadratic differentials is Strebel's book [33].

2.1. Quadratic differentials. Let S be a Riemann surface, and let ω_S denote its holomorphic cotangent bundle. A *meromorphic quadratic differential* ϕ on S is a meromorphic section of the line bundle $\omega_S^{\otimes 2}$. Two such differentials ϕ_1, ϕ_2 on surfaces S_1, S_2 are said to be equivalent if there is a biholomorphism $f: S_1 \rightarrow S_2$ such that $f^*(\phi_2) = \phi_1$.

In terms of a local co-ordinate z on S we can write a quadratic differential ϕ as

$$\phi(z) = \varphi(z) dz \otimes dz$$

with $\varphi(z)$ a meromorphic function. We write $\text{Zer}(\phi), \text{Pol}(\phi) \subset S$ for the subsets of zeroes and poles of ϕ respectively. The subset $\text{Crit}(\phi) = \text{Zer}(\phi) \cup \text{Pol}(\phi)$ is the set of *critical points* of ϕ .

At a point of $S \setminus \text{Crit}(\phi)$ there is a distinguished local co-ordinate w , uniquely defined up to transformations of the form $w \mapsto \pm w + \text{constant}$, with respect to which

$$\phi(w) = dw \otimes dw.$$

In terms of an arbitrary local co-ordinate z we have $w = \int \sqrt{\varphi(z)} dz$.

A quadratic differential ϕ determines two structures on $S \setminus \text{Crit}(\phi)$, namely a flat metric (called the ϕ -metric) and a foliation (the horizontal foliation). The ϕ -metric is defined locally by pulling back the Euclidean metric on \mathbb{C} using a distinguished co-ordinate w .

The horizontal foliation is given in terms of a distinguished co-ordinate by the lines $\text{Im}(w) = \text{constant}$.

The ϕ -metric and the horizontal foliation on $S \setminus \text{Crit}(\phi)$ together determine both the complex structure on S and the differential ϕ . Note that the set of quadratic differentials on a fixed surface S has a natural S^1 -action given by scalar multiplication : $\phi \mapsto e^{i\pi\theta} \cdot \phi$. This action has no effect on the ϕ -metric, but alters which in the circle of foliations defined by $\text{Im}(w/e^{i\pi\theta}) = \text{constant}$ is regarded as being horizontal.

In terms of a local co-ordinate z on S , the length of a smooth path γ in the ϕ -metric is

$$(2.1) \quad \ell_\phi(\gamma) = \int_\gamma |\varphi(z)|^{1/2} |dz|.$$

It is important to divide the critical set into a disjoint union

$$\text{Crit}(\phi) = \text{Crit}_{<\infty}(\phi) \cup \text{Crit}_\infty(\phi),$$

where $\text{Crit}_{<\infty}(\phi)$ consists of *finite critical points*, namely zeroes and simple poles, and $\text{Crit}_\infty(\phi)$ consists of *infinite critical points*, that is poles of order ≥ 2 . We write

$$S^\circ = S \setminus \text{Crit}_\infty(\phi)$$

for the complement of the infinite critical points.

Note that the integral (2.1) is well-defined for curves passing through points of $\text{Crit}_{<\infty}(\phi)$. This gives the surface S° the structure of a metric space, in which the distance between two points $p, q \in S^\circ$ is the infimum of the lengths of smooth curves in S° connecting p to q . The topology on S° defined by this metric agrees with the standard one induced from the surface S .

2.2. GMN differentials. All the quadratic differentials considered in this paper live on compact surfaces and have simple zeroes and at least one pole. Since it will be convenient to eliminate certain degenerate situations we make the following definition.

Definition 2.1. *A GMN differential is a meromorphic quadratic differential ϕ on a compact, connected Riemann surface S such that*

- (a) ϕ has simple zeroes,
- (b) ϕ has at least one pole,
- (c) ϕ has at least one finite critical point.

Condition (c) excludes polar types (2, 2) and (4) in genus 0; differentials of these types have unusual trajectory structures, and infinite automorphism groups.

Given a GMN differential (S, ϕ) we write g for the genus of the surface S and d for the number of poles of ϕ . The *polar type* of ϕ is the unordered collection of d integers

$m = \{m_i\}$ giving the orders of the poles of ϕ . We define

$$(2.2) \quad n = 6g - 6 + \sum_{i=1}^d (m_i + 1),$$

A GMN differential (S, ϕ) is said to be *complete* if ϕ has no simple poles, or in other words, if all $m_i \geq 2$. This is exactly the case in which the ϕ -metric on $S \setminus \text{Pol}(\phi)$ is complete. At the opposite extreme, (S, ϕ) is said to have *finite area* if ϕ has only simple poles, or in other words, if all $m_i = 1$.

2.3. Spectral cover and periods. Suppose that ϕ is a GMN differential on a compact Riemann surface S , with poles of order m_i at points $p_i \in S$. We can alternatively view ϕ as a holomorphic section

$$(2.3) \quad \varphi \in H^0(S, \omega_S(E)^{\otimes 2}), \quad E = \sum_i \left\lceil \frac{m_i}{2} \right\rceil \cdot p_i,$$

with simple zeroes at both the zeroes and the odd order poles of ϕ . The *spectral cover* of S defined by ϕ is the compact Riemann surface

$$\hat{S} = \{(p, l(p)) : p \in S, l(p) \in L_p \text{ such that } l(p) \otimes l(p) = \varphi(p)\} \subset L,$$

where L is the total space of the line bundle $\omega_S(E)$. This is a manifold because φ has simple zeroes.

The obvious projection map $\pi: \hat{S} \rightarrow S$ is a double cover, branched precisely over the zeroes and the odd order poles of the original meromorphic differential ϕ . There is a covering involution $\tau: \hat{S} \rightarrow \hat{S}$, commuting with the map π . The surface \hat{S} is connected because Definition 2.1 implies that π has at least one branch point.

We define the hat-homology group of the differential ϕ to be

$$\hat{H}(\phi) = H_1(\hat{S}^\circ; \mathbb{Z})^-,$$

where $\hat{S}^\circ = \pi^{-1}(S^\circ)$, and the superscript denotes the anti-invariant part for the action of the covering involution τ .

Lemma 2.2. *The group $\hat{H}(\phi)$ is free of rank n given by (2.2).*

Proof. The Riemann-Hurwitz formula applied to the spectral cover $\pi: \hat{S} \rightarrow S$ implies that

$$(2.4) \quad 2\hat{g} - 2 = 2(2g - 2) + (4g - 4 + \sum_{i=1}^d m_i) + (d - e),$$

where \hat{g} is the genus of \hat{S} , and e is the number of even m_i . The group $H_1(S^\circ; \mathbb{Z})$ is free of rank $2g + d - s - 1$, where s is the number of simple poles. Similarly, using

equation (2.4), and noting that each even order pole has two inverse images in \hat{S} , the group $H_1(\hat{S}^\circ; \mathbb{Z})$ is free of rank

$$r = 2\hat{g} + d + e - s - 1 = 8g - 6 + \sum_{i=1}^d m_i + 2d - s - 1.$$

Since the invariant part of $H_1(\hat{S}^\circ; \mathbb{Z})$ can be identified with $H_1(S^\circ; \mathbb{Z})$, the anti-invariant part $H_1(\hat{S}^\circ; \mathbb{Z})^-$ is therefore free of rank n . \square

The spectral cover \hat{S} comes equipped with a tautological section ψ of the line bundle $\pi^*(\omega_S(E))$ satisfying $\pi^*(\varphi) = \psi \otimes \psi$ and $\tau^*(\psi) = -\psi$. There is a canonical map $\eta: \pi^*(\omega_S) \rightarrow \omega_{\hat{S}}$ and we can form the composition

$$\mathcal{O}_{\hat{S}} \xrightarrow{\psi} \pi^*(\omega_S(E)) \xrightarrow{\eta(\hat{E})} \omega_{\hat{S}}(\hat{E}),$$

where $\hat{E} = \pi^{-1}(E)$. This defines a meromorphic 1-form on \hat{S} , which we also denote by ψ .

Since the canonical map η vanishes at the branch-points of π , the differential ψ is regular at the inverse images of the simple poles of ϕ , and hence restricts to a holomorphic 1-form on the open subsurface \hat{S}° . By construction ψ is anti-invariant for the action of the covering involution τ , and therefore defines a de Rham cohomology class

$$[\psi] \in H^1(\hat{S}^\circ; \mathbb{C})^-$$

called the period of ϕ . We choose to view this instead as a group homomorphism

$$Z_\phi: \hat{H}(\phi) \rightarrow \mathbb{C}.$$

2.4. Intersection forms. Consider a GMN differential ϕ on a Riemann surface S , and its spectral cover $\pi: \hat{S} \rightarrow S$. Write

$$\hat{D}_\infty = \pi^{-1}(\text{Crit}_\infty(\phi)).$$

Thus $\hat{S}^\circ = \hat{S} \setminus \hat{D}_\infty$. There are canonical maps of homology groups

$$H_1(\hat{S}^\circ; \mathbb{Z}) = H_1(\hat{S} \setminus \hat{D}_\infty; \mathbb{Z}) \xrightarrow{g} H_1(\hat{S}; \mathbb{Z}) \xrightarrow{h} H_1(\hat{S}, \hat{D}_\infty; \mathbb{Z}).$$

The intersection form on $H_1(\hat{S}; \mathbb{Z})$ is a non-degenerate, skew-symmetric pairing, and induces a degenerate skew-symmetric form

$$H_1(\hat{S}^\circ; \mathbb{Z}) \times H_1(\hat{S}^\circ; \mathbb{Z}) \rightarrow \mathbb{Z},$$

which we also call the intersection form, and write as $(\alpha, \beta) \mapsto \alpha \cdot \beta$. On the other hand, Lefschetz duality gives a non-degenerate pairing

$$(2.5) \quad \langle -, - \rangle: H_1(\hat{S} \setminus \hat{D}_\infty; \mathbb{Z}) \times H_1(\hat{S}, \hat{D}_\infty; \mathbb{Z}) \rightarrow \mathbb{Z}.$$

These bilinear forms restrict to the anti-invariant eigenspaces for the actions of the covering involutions.

For each pole $p \in S$ of ϕ of even order there is an associated *residue class*

$$\beta_p \in H_1(\hat{S}^\circ; \mathbb{Z})^-,$$

well-defined up to sign. It is obtained by taking the inverse image under π of a small loop in S° encircling the point p , and then orienting the two connected components so that the resulting class is anti-invariant.

The *residue* of ϕ at p is defined to be

$$(2.6) \quad \text{Res}_p(\phi) = Z_\phi(\beta_p) = \pm 2 \int_{\delta_p} \sqrt{\phi},$$

and is well-defined up to sign.

Lemma 2.3. *The classes $\beta_p \in H_1(\hat{S}^\circ; \mathbb{Z})^-$ are a \mathbb{Q} -basis for the kernel of the intersection form.*

Proof. If $p \in S$ is an even order pole of ϕ , let $\{s_p, t_p\}$ be the classes in $H_1(\hat{S}^\circ; \mathbb{Z})$ defined by small clockwise loops around the two inverse images of p in the spectral cover \hat{S} . Similarly, if $p \in S$ is a pole of odd order ≥ 3 , let $u_p \in H_1(\hat{S}^\circ; \mathbb{Z})$ be the class defined by a small loop around the single inverse image of p . Standard topology of surfaces shows that there is an exact sequence

$$0 \longrightarrow \mathbb{Z} \xrightarrow{i} \mathbb{Z}^{\oplus k} \xrightarrow{f} H_1(\hat{S}^\circ; \mathbb{Z}) \xrightarrow{g} H_1(\hat{S}; \mathbb{Z}) \longrightarrow 0,$$

where the map g is induced by the inclusion $\hat{S}^\circ \subset \hat{S}$, the map f sends the generators to the classes s_p, t_p and u_p respectively, and the image of i is the element $(1, 1, \dots, 1)$.

The covering involution exchanges s_p and t_p , and fixes u_p , and we have $\beta_p = \pm(s_p - t_p)$. Since the image of the map i lies in the invariant part of $H_1(\hat{S}; \mathbb{Z})$, the elements β_p are linearly independent. The intersection form on $H_1(\hat{S}; \mathbb{Z})^-$ is non-degenerate, so the kernel of the induced form on $H_1(\hat{S}^\circ; \mathbb{Z})^-$ is precisely the kernel of the surjective map

$$g^- : H_1(\hat{S}^\circ; \mathbb{Z})^- \rightarrow H_1(\hat{S}; \mathbb{Z})^-.$$

The group $H_1(\hat{S}; \mathbb{Z})^-$ has rank $2(\hat{g} - g)$, which by (2.4) is equal to $n - e$, where e is the number of even order poles of ϕ . Thus the kernel of g^- is spanned over \mathbb{Q} by the e elements β_p . \square

2.5. Moduli spaces. We now consider moduli spaces of GMN differentials of fixed polar type. For this purpose we fix a genus $g \geq 0$ and an unordered collection of $d \geq 1$ positive integers $m = \{m_i\}$.

Define $\text{Quad}(g, m)$ to be the set of equivalence-classes of pairs (S, ϕ) consisting of a compact, connected Riemann surface S of genus g , equipped with a GMN differential ϕ having polar type $m = \{m_i\}$.

Proposition 2.4. *The space $\text{Quad}(g, m)$ is either empty, or is a complex orbifold of dimension n given by (2.2).*

Proof. Let $\mathcal{M}(g, d)$ be the moduli stack of compact Riemann surfaces of genus g with an ordered set of d marked points (p_i) . This is a smooth algebraic stack of finite type over \mathbb{C} . Choose an ordering of the integers m_i , and let $\text{Sym}(m) \subset \text{Sym}(d)$ be the subgroup of the symmetric group consisting of permutations σ such that $m_{\sigma(i)} = m_i$.

At each point of $\mathcal{M}(g, d)/\text{Sym}(m)$ there is a Riemann surface S with a well-defined divisor $D = \sum_i m_i p_i$. The spaces of global sections $H^0(S, \omega_S^{\otimes 2}(D))$ fit together to form a vector bundle

$$(2.7) \quad \mathcal{H}(g, m) \longrightarrow \mathcal{M}(g, d)/\text{Sym}(m).$$

To see this, note first that if $g = 0$ then we can assume that the divisor D has degree at least 4, since otherwise the vector spaces are all zero, and the space $\text{Quad}(g, m)$ is empty. Serre duality therefore gives

$$H^1(S, \omega_S^{\otimes 2}(D)) \cong H^0(S, \omega_S(D)^\vee)^* = 0$$

which proves the claim. It then follows using Riemann-Roch that the rank of the bundle (2.7) is $3g - 3 + \sum_{i=1}^d m_i$.

The stack $\text{Quad}(g, m)$ is the Zariski open subset of $\mathcal{H}(g, m)$ consisting of sections with simple zeroes disjoint from the points p_i . Since $\mathcal{M}(g, d)$ is connected of dimension $3g - 3 + d$, the stack $\text{Quad}(g, m)$ is either empty, or is smooth and connected of dimension n .

The final step is to show that the automorphism groups of the relevant quadratic differentials are finite. This claim is clear if $g \geq 1$ or $d \geq 3$, because the same property holds for $\mathcal{M}(g, d)$. When $g = 0$ the claim is also clear if the total number of critical points is ≥ 3 . Since there is at least one pole, and the number of zeroes is $\sum m_i - 4$, the only other possibilities are polar types $(1, 3)$, (4) , (5) and $(2, 2)$.

In the first three of these cases there is a single quadratic differential up to equivalence, namely $\phi = z^k dz^{\otimes 2}$ with $k = -1, 0, 1$ respectively. The corresponding automorphism groups are $\{1\}$, $\mathbb{Z}_2 \ltimes \mathbb{C}^*$ and \mathbb{Z}_3 respectively. In the remaining case $(2, 2)$ the possible differentials are $\phi = r dz^{\otimes 2}/z^2$ for $r \in \mathbb{C}^*$. Each of these differentials has automorphism group $\mathbb{Z}_2 \ltimes \mathbb{C}^*$. By Definition 2.1(c), a GMN differential must have a zero or a simple pole; this exactly excludes the troublesome cases $(2, 2)$ and (4) . \square

Example 2.5. *Consider the case $g = 1, m = (1)$. The corresponding space $\text{Quad}(g, m)$ is empty, even though the expected dimension is $n = 2$. Indeed, this space parameterizes*

pairs (S, ϕ) , where S is a Riemann surface of genus 1, and ϕ is a meromorphic differential on S having only a simple pole. On the surface S the bundle ω_S is trivial, so ϕ defines a meromorphic function with a single simple pole. The Riemann-Roch theorem shows that no such function exists.

We shall often abuse notation by referring to the points of the space $\text{Quad}(g, m)$ as GMN differentials, and by denoting such a point simply by $\phi \in \text{Quad}(g, m)$. This is shorthand for the statement that ϕ is a GMN differential on a compact Riemann surface S , such that the equivalence class of the pair (S, ϕ) defines a point of the space $\text{Quad}(g, m)$.

The homology groups $H_1(S^\circ; \mathbb{Z})^-$ form a local system over the orbifold $\text{Quad}(g, m)$ because we can realise the spectral cover construction in families, and the Gauss-Manin connection gives a flat connection in the resulting bundle of anti-invariant homology groups. Often in what follows we will be studying a small neighbourhood

$$\phi_0 \in U \subset \text{Quad}(g, m)$$

of a fixed differential ϕ_0 . Whenever we do this we will tacitly assume that U is contractible, and use the Gauss-Manin connection to identify the hat-homology groups of all differentials in U .

2.6. Framings and the period map. As in the last section, we fix a genus $g \geq 0$ and a collection of $d \geq 1$ positive integers $m = \{m_i\}$. Let us also fix a free abelian group Γ of rank n given by (2.2).

As before, we consider pairs (S, ϕ) consisting of a Riemann surface S of genus g , equipped with a GMN differential ϕ of polar type $m = \{m_i\}$. A Γ -framing of such a pair (S, ϕ) is an isomorphism of groups

$$\theta: \Gamma \rightarrow \hat{H}(\phi).$$

Suppose (S_i, ϕ_i) for $i = 1, 2$ are two quadratic differentials as above, and $f: S_1 \rightarrow S_2$ is an isomorphism such that $f^*(\phi_2) = \phi_1$. Then f lifts to an isomorphism $\hat{f}: \hat{S}_1^\circ \rightarrow \hat{S}_2^\circ$, which is unique if we insist that it also satisfies $\hat{f}^*(\psi_2) = \psi_1$, where ψ_i are the distinguished 1-forms defined in Section 2.3.

Let $\text{Quad}^\Gamma(g, m)$ be the set of equivalence classes of triples (S, ϕ, θ) consisting of a compact, connected Riemann surface S of genus g equipped with a GMN differential ϕ of polar type $m = \{m_i\}$ together with a Γ -framing θ . We define triples (S_i, ϕ_i, θ_i) to be equivalent if there is an isomorphism $f: S_1 \rightarrow S_2$ such that $f^*(\phi_2) = \phi_1$ and such that

the distinguished lift \hat{f} makes the following diagram commute

$$(2.8) \quad \begin{array}{ccc} & \Gamma & \\ \theta_1 \swarrow & & \searrow \theta_2 \\ \hat{H}(\phi_1) & \xrightarrow{\hat{f}_*} & \hat{H}(\phi_2) \end{array}$$

We can define families of framed differentials in the obvious way, and the forgetful map

$$(2.9) \quad \text{Quad}^\Gamma(g, m) \rightarrow \text{Quad}(g, m)$$

is then an unbranched cover. Thus the set $\text{Quad}^\Gamma(g, m)$ is naturally a complex orbifold. The group $\text{Aut}(\Gamma)$ of automorphisms of the group Γ acts on $\text{Quad}^\Gamma(g, m)$, and the quotient orbifold is precisely $\text{Quad}(g, m)$. Note that $\text{Quad}^\Gamma(g, m)$ will not usually be connected, because the monodromy of the local system of hat-homology groups preserves the intersection form, and hence cannot relate all different framings of a given differential. But since all such framings *are* related by the action of $\text{Aut}(\Gamma)$, the different connected components of $\text{Quad}^\Gamma(g, m)$ are all isomorphic.

The period of a framed GMN differential (S, ϕ, θ) can be viewed as a map $Z_\phi \circ \theta: \Gamma \rightarrow \mathbb{C}$. This gives a well-defined *period map*

$$(2.10) \quad \pi: \text{Quad}^\Gamma(g, m) \rightarrow \text{Hom}_{\mathbb{Z}}(\Gamma, \mathbb{C}).$$

In Section 4.7 we shall prove that, with the exception of the six special cases considered in the next section, the space $\text{Quad}^\Gamma(g, m)$ is a complex manifold, and the period map π is a local homeomorphism.

2.7. Generic automorphisms. In certain special cases the orbifolds $\text{Quad}(g, m)$ and $\text{Quad}^\Gamma(g, m)$ have non-trivial generic automorphism groups. In this section we classify the polar types when this occurs.

Lemma 2.6. *The generic automorphism group of a point of $\text{Quad}(g, m)$ is trivial, with the exception of the polar types*

$$(5); \quad (6); \quad (1, 1, 2); \quad (3, 3); \quad (1, 1, 1, 1),$$

in genus $g = 0$, and the polar type $m = (2)$ in genus $g = 1$.

Proof. Suppose first that if $g = 0$ then $d \geq 5$, and that if $g = 1$ then $d \geq 2$. With these assumptions it is well-known that the stack $\mathcal{M}(g, d)/\text{Sym}(d)$ parameterizing compact Riemann surfaces of genus g with an unordered collection of d marked points has trivial generic automorphism group. The same is therefore true of the stack $\mathcal{M}(g, d)/\text{Sym}(m)$ appearing in the proof of Proposition 2.4. The space $\text{Quad}(g, m)$ is an open subset of a vector bundle over this stack, so again, the generic automorphism group is trivial.

Consider the case $g = 1$ and $d = 1$. The stack $\text{Quad}(g, m)$ then parameterizes pairs consisting of a Riemann surface S of genus 1, together with a meromorphic function

on S with simple zeroes and a single pole, necessarily of order $m \geq 2$. For a generic such surface S , the group of automorphisms preserving the pole is generated by a single involution, and using Riemann-Roch it is easily seen that if $m \geq 3$ then the zeroes of the generic such function are not permuted by this involution.

When $g = 0$ Riemann-Roch shows that there exist differentials with any given configuration of zeroes and poles, providing only that the number k of zeroes is equal to $\sum m_i - 4$. Thus if a generic point $\phi \in \text{Quad}(g, m)$ has non-trivial automorphisms, then $|\text{Crit}(\phi)| \leq 4$. Moreover, if $|\text{Crit}(\phi)| = 4$ then the critical points must consist of two pairs of the same type, since the generic automorphism group of $\mathcal{M}(0, 4)/\text{Sym}(4)$ acts on the marked points via permutations of type $(ab)(cd)$. Similarly, if $|\text{Crit}(\phi)| = 3$ then at least two of the critical points must be of the same type.

Suppose that the generic point of $\text{Quad}(0, m)$ does have non-trivial automorphisms. Since there is at least one pole, we must have $0 \leq k \leq 3$. We cannot have $k = 3$ since there would then be 4 critical points whose types do not match in pairs. If $k = 2$ there must be two poles of the same degree, giving the $(3, 3)$ case, or a single pole, giving the (6) case. If $k = 1$ there must be just one pole, which gives the case (5) , since if there were 2 poles they would have to have the same degree. Finally, if $k = 0$ we get the cases $(1, 1, 2)$ and $(1, 1, 1, 1)$, since the cases $(2, 2)$ and (4) have already been excluded by the definition of a GMN differential, and the case $(1, 3)$ leads to a single differential with trivial automorphism group, as discussed in the proof of Proposition 2.4. \square

Examples 2.7. *We consider differentials $(S, \phi) \in \text{Quad}(g, m)$ corresponding to some of the exceptional cases in the statement of Lemma 2.6.*

- (a) *Consider the case $g = 0$ and $m = (1, 1, 2)$. Taking the simple poles to be at $\{0, \infty\} \in \mathbb{P}^1$ we can write any such differential in the form*

$$\phi(z) = \frac{c dz^{\otimes 2}}{z(z-1)^2}$$

for some $c \in \mathbb{C}^$. Thus ϕ is invariant under the automorphism $z \mapsto 1/z$. The spectral cover \hat{S} is again \mathbb{P}^1 with co-ordinate $w = \sqrt{z}$ and covering involution $w \mapsto -w$. The automorphism $z \mapsto 1/z$ lifts to the automorphism $w \mapsto 1/w$ of the open subsurface $\hat{S}^\circ = \mathbb{P}^1 \setminus \{\pm 1\}$ and acts trivially on the hat-homology group, which is $H_1(\hat{S}^\circ; \mathbb{Z}) = \mathbb{Z}$. Thus every element of $\text{Quad}^\Gamma(g, m)$ has automorphism group \mathbb{Z}_2 .*

- (b) *Consider the case $g = 0$, $m = (3, 3)$. Any such differential is of the form*

$$\phi(z) = (tz + 2s + tz^{-1}) \frac{dz^{\otimes 2}}{z^2},$$

for constants $s, t \in \mathbb{C}^$ with $s \pm t \neq 0$, and is invariant under $z \mapsto 1/z$. The spectral cover \hat{S} has genus 1, and the covering involution is the hyperelliptic involution. The open subset \hat{S}° is the complement of 2 points, the inverse images*

of the poles of ϕ . The automorphism $z \mapsto 1/z$ of \mathbb{P}^1 lifts to a translation by a 2-torsion point of \hat{S} . It acts trivially on the hat-homology group, which is $H_1(\hat{S}; \mathbb{Z}) = \mathbb{Z}^{\oplus 2}$. Thus every point of $\text{Quad}^\Gamma(g, m)$ has automorphism group \mathbb{Z}_2 .

(c) Consider the case $g = 0$, $m = (1, 1, 1, 1)$. Such differentials are of the form

$$\phi(z) = \frac{dz^{\otimes 2}}{p_4(z)},$$

where $p_4(z)$ is a monic polynomial of degree 4 with distinct roots, and are invariant under any automorphism of \mathbb{P}^1 permuting these roots. The spectral cover \hat{S} has genus 1, and the covering involution is the hyper-elliptic involution. The automorphisms of \mathbb{P}^1 preserving ϕ lift to translations by 2-torsion points of \hat{S} . These automorphisms act trivially on the hat-homology group, which is $H_1(\hat{S}; \mathbb{Z}) = \mathbb{Z}^{\oplus 2}$. Thus every point of $\text{Quad}^\Gamma(g, m)$ has automorphism group $\mathbb{Z}_2^{\oplus 2}$.

In each of the other cases of Lemma 2.6 the orbifold $\text{Quad}^\Gamma(g, m)$ also has non-trivial generic automorphism group. The case $g = 0$, $m = (5)$ is elementary, and the case $g = 0$, $m = (6)$ is very similar to Example 2.7(a). The case $g = 1$, $m = (2)$ is treated in Example 12.4 below.

3. TRAJECTORIES AND GEODESICS

In this section we focus on the global trajectory structure of a fixed quadratic differential, and the basic properties of the geodesic arcs of the associated flat metric. This material is all well-known, but since it forms the basis for much of what follows we thought it worthwhile to give a fairly detailed treatment. The reader can find proofs and further explanations in Strebel's book [33].

3.1. Trajectories. Let ϕ be a meromorphic quadratic differential on a compact Riemann surface S . A *straight arc* in S is a smooth path $\gamma: I \rightarrow S \setminus \text{Crit}(\phi)$, defined on an open interval $I \subset \mathbf{R}$, which makes a constant angle $\pi\theta$ with the horizontal foliation. In terms of a distinguished local co-ordinate w as in Section 2.1 the condition is that the function $\text{Im}(w/e^{i\pi\theta})$ should be constant along γ . The phase θ of a straight arc is a well-defined element of \mathbf{R}/\mathbb{Z} ; in the case $\theta = 0$ the arc is said to be *horizontal*.

We make the convention that all straight arcs are parameterized by arc-length in the ϕ -metric. Straight arcs differing by a reparameterization (necessarily of the form $t \mapsto \pm t + \text{constant}$) will be regarded as being the same. A straight arc is called *maximal* if it is not the restriction of a straight arc defined on a larger interval. A maximal horizontal straight arc is called a *trajectory*. Every point of $S \setminus \text{Crit}(\phi)$ lies on a unique trajectory, and any two trajectories are either disjoint or coincide.

We define a *saddle trajectory* to be a trajectory γ whose domain of definition is a finite interval $(a, b) \subset \mathbf{R}$. Since S is compact, we can then extend γ to a continuous path $\gamma: [a, b] \rightarrow S$, whose endpoints $\gamma(a)$ and $\gamma(b)$ are finite critical points of ϕ . We tend not to distinguish between the saddle trajectory γ and its closure. By a *closed saddle trajectory* we mean a saddle trajectory whose endpoints coincide.

More generally, a *saddle connection* is a maximal straight arc of some phase θ whose domain of definition is a finite interval. Thus a saddle trajectory is a horizontal saddle connection, and a saddle connection of phase θ is a saddle trajectory for the rotated differential $e^{-i\pi\theta} \cdot \phi$.

If a trajectory γ intersects itself, then it must be periodic, and have domain $I = \mathbf{R}$. In this situation we usually restrict the domain of γ to a primitive period $[a, b] \subset \mathbf{R}$, and refer to γ as a *closed trajectory*. By a *finite-length trajectory* we mean either a closed trajectory or a saddle trajectory.

3.2. Hat-homology classes. Let us again fix a meromorphic quadratic differential ϕ on a compact Riemann surface S . The inverse image of the horizontal foliation of $S \setminus \text{Crit}(\phi)$ under the covering map π defines a horizontal foliation on $\hat{S} \setminus \pi^{-1} \text{Crit}(\phi)$. In more detail, the 1-form ψ of Section 2.3 can be written locally as $\psi = d\hat{w}$, and the horizontal foliation of \hat{S} is then given by the lines $\text{Im}(\hat{w}) = \text{constant}$. This foliation can be canonically oriented by insisting that ψ evaluated on the tangent vector to the oriented foliation should lie in $\mathbf{R}_{>0}$ rather than $\mathbf{R}_{<0}$. Note that since ψ is anti-invariant, the covering involution τ preverses the horizontal foliation on \hat{S} , but reverses its orientation.

Suppose that $\gamma: [a, b] \rightarrow S$ is a finite-length trajectory. The inverse image $\pi^{-1}(\gamma)$ is then a closed curve in the spectral cover \hat{S} , which could be disconnected (if γ is a closed trajectory), or singular (if γ is a closed saddle trajectory, see Figure 7). In all cases we orient $\pi^{-1}(\gamma)$ according to the orientation discussed in the previous paragraph. Since the covering involution flips this orientation, we obtain a class $\hat{\gamma} \in \hat{H}(\phi)$ called the *hat-homology class*¹ of the trajectory γ . Note that, by definition, it satisfies $Z_\phi(\hat{\gamma}) \in \mathbf{R}_{>0}$.

Similar remarks apply to maximal straight arcs of finite-length and nonzero phase θ . The only difference is that we orient the inverse image of the arcs on \hat{S} by insisting that ψ evaluated on the tangent vector should have positive imaginary part. This means that the corresponding hat-homology classes have periods $Z_\phi(\hat{\gamma})$ lying in the upper half-plane.

¹With this definition it is not necessarily the case that $\hat{\gamma}$ is primitive, cf. Figure 7. In the literature one often sees a more complicated definition of the hat-homology class of a saddle trajectory which boils down to taking the unique primitive multiple of our $\hat{\gamma}$.




FIGURE 7. A closed saddle trajectory γ , and its preimages γ^\pm in the spectral cover, whose union define its (imprimitive) hat-homology class.

3.3. Critical points. We now describe the local structure of the horizontal foliation near a critical point of a meromorphic quadratic differential, following Strebel [33, §6].

Let ϕ be a meromorphic quadratic differential on a Riemann surface S . Suppose first that $p \in \text{Crit}_{<\infty}(\phi)$ is either a simple pole of ϕ , in which case we set $k = -1$, or a zero of some order $k \geq 1$. Then there are local co-ordinates t such that

$$\phi(t) = c^2 \cdot t^k dt \otimes dt, \quad c = \frac{1}{2}(k+2).$$

At nearby points of $S \setminus \{p\}$, a distinguished local co-ordinate is $w = t^{\frac{1}{2}(k+2)}$. The local trajectory structure is illustrated in the cases $k = \pm 1$ in Figure 8.




FIGURE 8. Local trajectory structures at a simple zero and a simple pole.

Note that three horizontal rays emanate from each simple zero; this trivalent structure will be the basic reason for the link with triangulations.

Next suppose that $p \in \text{Crit}_\infty(\phi)$ is a pole of order 2. Then there are local co-ordinates t such that

$$\phi(t) = r \frac{dt^{\otimes 2}}{t^2},$$

for some well-defined constant $r \in \mathbb{C}^*$. The residue of ϕ at p is

$$(3.1) \quad Z_\phi(\beta_p) = \text{Res}_p(\phi) = \pm 4\pi i \sqrt{r},$$

and is well-defined up to sign.

At nearby points of $S \setminus \{p\}$ any branch of the function $w = \sqrt{r} \log(t)$ is a distinguished local co-ordinate, and the structure of the horizontal foliation near p is determined by the residue as follows:

- (i) if $\text{Res}_p(\phi) \in \mathbb{R}$ the foliation is by concentric circles centred on the pole;
- (ii) if $\text{Res}_p(\phi) \in i\mathbb{R}$ the foliation is by radial arcs emanating from the double pole;
- (iii) if $\text{Res}_p(\phi) \notin \mathbb{R} \cup i\mathbb{R}$ the leaves of the foliation are logarithmic spirals which wrap onto the pole.




FIGURE 9. Local trajectory structures at a double pole.

These three cases are illustrated in Figure 9. In cases (ii) and (iii) there is a neighbourhood $p \in U \subset S$ such that any trajectory entering U tends to p .

Finally, suppose that $p \in \text{Crit}_\infty(\phi)$ is a pole of order $m > 2$. If m is odd, there are local co-ordinates t such that

$$\phi(t) = c^2 \cdot t^{-m} dt \otimes dt, \quad c = \frac{1}{2}(2 - m).$$


as before. If $m \geq 4$ is even, there are local co-ordinates t such that

$$\phi(t) = \left(ct^{-m/2} + \frac{b}{t}\right)^2 dt \otimes dt, \quad c = \frac{1}{2}(2 - m).$$

The residue of ϕ at p is then

$$Z_\phi(\beta_p) = \text{Res}_p(\phi) = \pm 4\pi ib,$$

and is well-defined up to sign.


 FIGURE 10. Local trajectory structures at poles of order $m = 3, 4, 5$.

The trajectory structure in these cases is illustrated in Figure 10. There is a neighbourhood $p \in U \subset S$ and a collection of $m - 2$ distinguished tangent directions v_i at p , such that any trajectory entering U eventually tends to p and becomes asymptotic to one of the v_i .

3.4. Global trajectories. Let ϕ be a GMN differential on a compact Riemann surface S . We now consider the global structure of the horizontal foliation of ϕ , again following Strebel [33, §§ 9–11]. Every trajectory of ϕ falls into exactly one of the following categories:

- (1) *saddle trajectories* approach finite critical points at both ends;
- (2) *separating trajectories* approach critical points at each end, one finite and one infinite;
- (3) *generic trajectories* approach infinite critical points at both ends;
- (4) *closed trajectories* are simple closed curves in $S \setminus \text{Crit}(\phi)$;
- (5) *divergent trajectories* are recurrent in at least one direction.

Since only finitely many horizontal arcs emerge from each finite critical point, the number of saddle trajectories and separating trajectories is finite. Removing these from S , together with the critical points $\text{Crit}(\phi)$, the remaining open surface splits as a disjoint union of connected components which can be classified as follows²

- (1) A *half-plane* is equivalent to the upper half-plane

$$\{z \in \mathbb{C} : \text{Im}(z) > 0\} \subset \mathbb{C}$$

equipped with the differential $dz^{\otimes 2}$. It is swept out by generic trajectories which connect a fixed pole of order $m > 2$ to itself. The boundary is made up of saddle trajectories and separating trajectories.

- (2) A *horizontal strip* is equivalent to a region

$$\{z \in \mathbb{C} : a < \text{Im}(z) < b\} \subset \mathbb{C},$$

equipped with the differential $dz^{\otimes 2}$. It is swept out by generic trajectories connecting two (not necessarily distinct) poles of arbitrary order $m \geq 2$. Each component of the boundary is made up of saddle trajectories and separating trajectories.

- (3) A *ring domain* is equivalent to a region

$$\{z \in \mathbb{C} : a < |z| < b\} \subset \mathbb{C}^*,$$

equipped with the differential $r dz^{\otimes 2}/z^2$ for some $r \in \mathbf{R}_{<0}$. It is swept out by closed trajectories. Each component of the boundary is either made up of saddle trajectories or is a single double pole of ϕ with real residue.

²See [33, Section 11.4]. Strictly speaking the decomposition is into *maximal* horizontal strips, half-planes etc, but since all such domains we consider will be maximal, we drop the qualifier. Recall that we have outlawed various degenerate cases: by assumption ϕ has at least one finite critical point, and at least one pole.

- (4) A *spiral domain* is defined to be the interior of the closure of a divergent trajectory. The only fact we shall need is that the boundary of a spiral domain is made up of saddle trajectories. In particular there are no infinite critical points in the closure of a spiral domain.

A ring domain A will be called *degenerate* if one of its boundary components consists of a double pole p . The residue $\text{Res}_p(\phi)$ is then necessarily real, and A consists of closed trajectories encircling p . Conversely, any double pole p with real residue is contained in a degenerate ring domain. A ring domain A will be called *strongly non-degenerate* if its boundary consists of two, pairwise disjoint, simple closed curves on S . Not all non-degenerate ring domains are strongly non-degenerate; for example, in the case of finite area differentials, there is a dense subspace of $\text{Quad}(g, m)$ consisting of differentials which have a single dense ring domain [33, Theorem 25.2].

3.5. Saddle-free differentials. We say that a GMN differential is *saddle-free* if it has no saddle trajectories. The following simple but crucial observation comes from [14, §6.3].

Lemma 3.1. *If a GMN differential ϕ is saddle-free, and $\text{Crit}_\infty(\phi)$ is non-empty, then ϕ has no closed or divergent trajectories.*

Proof. Since $\text{Crit}_\infty(\phi)$ is non-empty the surface S cannot be the closure of a spiral domain. On the other hand, the boundary of a spiral domain consists of saddle trajectories. Thus there can be no spiral domains, and hence no divergent trajectories. Similarly the boundary of a ring domain must contain saddle trajectories, except for the case when both boundary components are double poles with real residue. This can only occur when $g = 0$ and the polar type is $m = (2, 2)$; such differentials are not GMN since they have no finite critical points. \square

Let ϕ be a saddle-free GMN differential such that $\text{Crit}_\infty(\phi)$ is non-empty. Removing the finitely many separating trajectories from $S \setminus \text{Crit}(\phi)$ gives an open surface which is a disjoint union of horizontal strips and half-planes swept out by generic trajectories.

Each of the two components of the boundary of a horizontal strip contains exactly one finite critical point of ϕ . If these are both zeroes, then embedded in the surface there are two possibilities, depending on whether the two zeroes are distinct or coincide; we call the corresponding strips *regular* or *degenerate* respectively. These two possibilities are illustrated in Figure 12; note though that the two double poles in the first of these pictures could well coincide on the surface.

A horizontal strip containing a simple pole in one of its boundary components is almost always of the form illustrated in Figure 13. The one exception occurs in genus 0 and polar type $(1, 1, 2)$: the moduli space of such differentials consists of a single \mathbb{C}^* -orbit,




FIGURE 11. The generic (dotted) and separating trajectories (solid) for a saddle-free GMN differential having only double poles. All horizontal strips in the picture are non-degenerate.




FIGURE 12. Two types of strip, regular and degenerate.




FIGURE 13. Horizontal strip with a simple pole on its boundary; the simple pole is in the centre of the diagram with a double pole above and a simple zero below.

and the trajectory structure for a generic element consists of a single horizontal strip containing two simple poles in its boundary.

3.6. Standard saddle connections. Let ϕ be a saddle-free GMN differential on a Riemann surface S , and assume that $\text{Crit}_\infty(\phi)$ is non-empty. The interior of each horizontal strip is equivalent to a strip in \mathbb{C} equipped with the differential $dz^{\otimes 2}$. In each

such strip h there is a unique saddle connection ℓ_h connecting the two finite critical points on the opposite sides of the strip.




FIGURE 14. A horizontal strip in \mathbb{C} with its standard saddle connection.

Since ϕ is saddle-free, ℓ_h must have nonzero phase. As in Section 3.1, there is an associated hat-homology class $\alpha_h \in \hat{H}(\phi)$, which by definition satisfies $\text{Im } Z_\phi(\alpha_h) > 0$. We call the arcs ℓ_h the *standard saddle connections* of the differential ϕ . The classes α_h will be called the *standard saddle classes*.

Lemma 3.2. *The standard saddle classes α_h form a basis for the group $\hat{H}(\phi)$.*

Proof. In each horizontal strip h_i we can choose a generic trajectory and then take one of its two lifts to the spectral cover to give a class δ_{h_i} in the relative homology group of (2.5). The intersection number $\langle \alpha_{h_i}, \delta_{h_j} \rangle$ is then nonzero precisely if $h_i = h_j$, in which case it is ± 1 . Thus the elements α_{h_i} are linearly independent. Lemma 2.2 states that the group $\hat{H}(\phi)$ is free of rank n given by equation (2.2). To complete the proof it will be enough to show that this is also the number of horizontal strips of ϕ .

By a transverse orientation of a separating trajectory we mean a continuous choice of normal direction; for each separating trajectory there are two possible choices. We orient the separating trajectories in the boundary of a horizontal strip by taking the inward pointing normal direction. Each horizontal strip then has four transversally oriented separating trajectories in its closure; for a degenerate strip, two of these consist of different orientations of the same trajectory. Similarly, each half-plane has two such oriented trajectories. Moreover, every oriented separating trajectory occurs as the boundary of exactly one half-plane or horizontal strip.

Let x be the number of horizontal strips, and s the number of simple poles. Three horizontal arcs emanate from each zero, and one from each simple pole, and each of these forms the end of a separating trajectory. Each pole of order $m \geq 3$ is surrounded by $m - 2$ half-planes, so the total number of these is $s + \sum_{i=1}^d (m_i - 2)$. Thus we get an equality

$$4x + 2s + 2 \sum_{i=1}^d m_i - 4d = 6(4g - 4 + \sum_{i=1}^d m_i) + 2s.$$

Simplifying this expression gives $x = n$. □

3.7. Geodesics. Let ϕ be a meromorphic quadratic differential on a Riemann surface S . Recall from Section 2.1 that ϕ induces a metric space structure on the open subsurface $S^\circ = S \setminus \text{Crit}_\infty(\phi)$. A ϕ -geodesic is defined to be a locally-rectifiable path $\gamma: [0, 1] \rightarrow S^\circ$ which is locally length-minimizing. Note that it is not assumed that γ is the shortest path between its endpoints.

It follows immediately from the definition of the ϕ -metric that any straight arc is a ϕ -geodesic, and that conversely, in a neighbourhood of a non-critical point of ϕ , any geodesic is a straight arc. Similarly, using the canonical co-ordinate systems of Section 3.3, it is easy to determine the local behaviour of geodesics near a finite critical point of ϕ . Here we briefly summarize the results of this analysis, and refer the reader to Strebel [33, §8] for more details.

In a neighbourhood of a zero p of ϕ of order k , any two points are joined by a unique geodesic, which is also the shortest curve in S° connecting these points. This unique geodesic is either a straight arc not passing through p , or is composed of two radial straight arcs emanating from p . This second situation occurs precisely if the angle between the radial arcs is $\geq 2\pi/(k+2)$.




FIGURE 15. Geodesic segments near a simple zero.

In a neighbourhood of a simple pole p of ϕ , any two points are connected by at least one geodesic, but uniqueness of geodesics fails: some pairs of points are connected by more than one straight arc. Moreover, a geodesic need not be the shortest path between its endpoints: it is length-minimizing locally, but not necessarily globally. Note however, that no geodesic contains the point p in its interior: the only geodesics passing through p begin or end there.

From these local descriptions, it immediately follows that any geodesic in S° is a union of (closures of) straight arcs, joined at zeroes of ϕ . In particular, any geodesic connecting points of $\text{Crit}_{<\infty}(\phi)$ is a union of saddle connections. Of course, the phases of the constituent saddle connections will usually be different.




FIGURE 16. Geodesic segments near a simple pole, and their inverse images under the square-root map. Note that the pulled-back differential has a non-critical point at the inverse image of the pole.

3.8. Gluing surfaces along geodesics. It will be useful in what follows to glue Riemann surfaces equipped with quadratic differentials along closed curves made up of unions of saddle connections. We will see some particular examples of this construction in Sections 5.5 and 6.4 below.

Consider a topological surface \mathbb{S} with boundary. By a quadratic differential on \mathbb{S} we simply mean a quadratic differential on the interior of \mathbb{S} , that is a quadratic differential on a Riemann surface whose underlying topological surface is the interior of \mathbb{S} . We say that two such surfaces \mathbb{S}_i equipped with differentials ϕ_i are equivalent if there is a homeomorphism $f: \mathbb{S}_1 \rightarrow \mathbb{S}_2$ which restricts to a biholomorphism on the interiors and satisfies $f^*(\phi_2) = \phi_1$.

Given an integer $k \geq 0$ we denote by $\mathbb{V}_k \subset \mathbb{C}$ the closed sector bounded by the rays of argument 0 and $2\pi(k+1)/(k+2)$. We equip the interior of \mathbb{V}_k with the differential

$$\phi_k(t) = c^2 \cdot t^k dt^{\otimes 2}, \quad c = \frac{1}{2}(k+2).$$

Thus, for example, $\mathbb{V}_0 \subset \mathbb{C}$ is the closed upper half-plane equipped with the standard differential $\phi_0(t) = dt^{\otimes 2}$ on its interior. In general the differential ϕ_k extends holomorphically over a neighbourhood of the boundary of \mathbb{V}_k , and when $k > 0$, the boundary $\partial\mathbb{V}_k$ then consists of two horizontal trajectories of ϕ_k meeting at a zero of order k .

Note that the map $z^2 = t^{k+2}$ gives an equivalence

$$(3.2) \quad (\mathbb{C} \setminus \mathbb{V}_k, \phi_k(t)) \cong (\mathfrak{h}, dz^{\otimes 2}).$$

Thus a copy of \mathbb{V}_k can be glued to a copy of \mathbb{V}_0 in such a way that the differentials ϕ_k and ϕ_0 on the interiors extend to a well-defined differential on \mathbb{C} .

If ϕ is a quadratic differential on a topological surface \mathbb{S} with boundary, we say that the pair (S, ϕ) has a *gluable boundary* if each point $x \in \partial\mathbb{S}$ has a neighbourhood which is equivalent to a neighbourhood of $0 \in \mathbb{V}_k$ for some $k \geq 0$. In particular it follows that the boundary $\partial\mathbb{S}$ is either a union of saddle trajectories or a single closed trajectory. Note, however, that the gluable boundary condition is a much stronger statement: if $z \in \partial\mathbb{S}$

is a zero of ϕ of order k , then there are $k + 2$ horizontal trajectories in \mathbb{S} emanating from z , two of which lie in the boundary.

Suppose that S is a Riemann surface equipped with a meromorphic differential ϕ having simple zeroes, and that $\gamma \subset S$ is a separating simple closed curve which is either a closed trajectory or a union of saddle trajectories. Cutting the underlying topological surface \mathbb{S} along γ we can view it as a union of two surfaces with boundary \mathbb{S}_{\pm} glued along the curve γ . The assumption that ϕ has simple zeroes then immediately implies that the pairs $(\mathbb{S}_{\pm}, \phi|_{\mathbb{S}_{\pm}})$ have glueable boundaries in the sense described above.

Conversely, suppose that \mathbb{S}_{\pm} are two smooth, oriented surfaces with boundary, each with a single boundary component $\partial\mathbb{S}_{\pm}$, and each equipped with quadratic differentials ϕ_{\pm} .

Lemma 3.3. *Suppose that the pairs $(\mathbb{S}_{\pm}, \phi_{\pm})$ have glueable boundaries, and that the ϕ_{\pm} -lengths of the boundaries $\partial\mathbb{S}_{\pm}$ are equal. Then there is a Riemann surface S whose underlying topological surface \mathbb{S} is obtained by gluing the surfaces \mathbb{S}_{\pm} along their boundaries, and a meromorphic differential ϕ on S which coincides with the differentials ϕ_{\pm} on the interiors of the two subsurfaces $\mathbb{S}_{\pm} \subset \mathbb{S}$.*

Proof. Parameterize the two boundary components $\partial\mathbb{S}_{\pm}$ by arc-length in the ϕ_{\pm} -metric, and then identify them. When we do this we have the freedom to choose the rotation of the two surfaces relative to each other, and we can therefore ensure that zeroes of ϕ_{\pm} do not become identified. The fact that the quadratic differentials ϕ_{\pm} glue together then follows from the equivalence (3.2). \square

Remarks 3.4. (a) *It is clear from the proof of Lemma 3.3 that the surface S is not uniquely determined by the pairs $(\mathbb{S}_{\pm}, \phi_{\pm})$: we can rotate the subsurfaces \mathbb{S}_{\pm} relative to one another.*

(b) *The glueable boundary assumption is necessary: one cannot always glue differentials on surfaces whose boundaries are made up of saddle trajectories. Indeed, otherwise one could take a degenerate ring domain whose boundary consists of $i \geq 1$ saddle trajectories, and glue it to itself to obtain a meromorphic differential on a sphere with 2 double poles and i simple zeroes. This cannot exist by Riemann-Roch.*

4. PERIOD CO-ORDINATES

The aim of this section is to prove that the period map (2.10) on the space of framed differentials is a local isomorphism. For finite area differentials this is standard, but for the more general meromorphic differentials considered here there does not seem to be a proof in the literature. The reader prepared to take this result on trust can skip to the next section. We begin by considering geodesics for the metric defined by a GMN

differential ϕ , and the way in which these change as ϕ moves in the corresponding space $\text{Quad}(g, m)$.

4.1. Existence and uniqueness of geodesics. Let ϕ be a meromorphic quadratic differential on a compact Riemann surface S . As in Section 2.1 we equip the open subsurface $S^\circ = S \setminus \text{Crit}(\phi)$ with the metric space structure induced by the ϕ -metric. In this section we state some well-known global existence and uniqueness properties for geodesics on this surface. A more detailed treatment can be found in [33, §§14–18].

Given points $p, q \in S^\circ$, we denote by $\mathcal{C}(p, q)$ the set of all rectifiable paths $\gamma: [0, 1] \rightarrow S^\circ$ connecting p to q . We equip this set with the topology of uniform convergence. Two curves in $\mathcal{C}(p, q)$ are considered homotopic if they are homotopic relative to their endpoints through paths in S° . We denote by $\ell_\phi(\gamma)$ the length of a curve $\gamma \in \mathcal{C}(p, q)$. A curve in $\mathcal{C}(p, q)$ will be called a *minimal geodesic* if no homotopic path has smaller length; any such curve is locally length-minimising, and hence a geodesic.

The following result is well-known.

Theorem 4.1. (a) *the subset of curves in $\mathcal{C}(p, q)$ representing a given homotopy class is open and closed,*
 (b) *the function sending a curve in $\mathcal{C}(p, q)$ to its length is lower semi-continuous,*
 (c) *for any $L > 0$, the subset of curves in $\mathcal{C}(p, q)$ of length $\leq L$ which are parameterized proportional to arc-length is compact,*
 (d) *every homotopy class of curves in $\mathcal{C}(p, q)$ contains at least one minimal geodesic,*
 (e) *if ϕ has no simple poles then geodesics in $\mathcal{C}(p, q)$ are homotopic only if they are equal.*

Proof. Since the surface S is assumed compact, the metric space S° is proper, which is to say that all closed, bounded subsets are compact. It is also clear that any two points of S° can be connected by a rectifiable path. The statements (a) - (d) hold for all metric spaces with these two properties: see for example [29, Section 1.4]. Part (e) is proved by Strebel [33, Theorem 16.2]. \square

If the differential ϕ has no simple poles, Theorem 4.1 implies that all geodesics are minimal. If ϕ has simple poles the situation is more complicated: a given homotopy class may contain more than one geodesic representative, and not all such representatives need be minimal.

Lemma 4.2. *For any $L > 0$, there are only finitely many geodesics $\gamma \in \mathcal{C}(p, q)$ with $\ell_\phi(\gamma) \leq L$.*

Proof. First assume that ϕ has no simple poles. It follows from Theorem 4.1(c) that the subset of $\mathcal{C}(p, q)$ consisting of curves of length $\leq L$ has only finitely many connected

components. In particular, by part (a), there can only be finitely many homotopy classes of such curves. But, by part (e), a geodesic is determined by its homotopy class, so the result follows. In the general case, take a covering $\pi: \tilde{S} \rightarrow S$ branched at all simple poles of ϕ , and consider the pulled-back differential $\tilde{\phi} = \pi^*(\phi)$. Any ϕ -geodesic in S can be lifted to a $\tilde{\phi}$ -geodesic in \tilde{S} of the same length. Since $\tilde{\phi}$ has no simple poles, this reduces us to the previous case. \square

4.2. Varying the differential. Our next step is to study the way geodesics of a GMN differential move as the differential varies in its moduli space. Fix a genus $g \geq 0$ and a collection of $d \geq 1$ positive integers $m = \{m_i\}$. Recall from the proof of Proposition 2.4 that, when it is non-empty, the space $\text{Quad}(g, m)$ is an open subset of a vector bundle

$$\mathcal{H}(g, m) \longrightarrow \mathcal{M}(g, d) / \text{Sym}(m).$$

The fibre of this bundle over a marked curve $(S, (p_i))$ is the space of global sections of the line bundle $\omega_S^{\otimes 2}(\sum_i m_i p_i)$.

Let us consider a fixed differential $\phi_0 \in \text{Quad}(g, m)$, which we view as a base-point, and consider an open ball³ $\phi_0 \in Q \subset \text{Quad}(g, m)$. By Ehresmann's theorem, the universal curve over $\mathcal{M}(g, d)$ pulls back to a differentiably locally-trivial fibre bundle over Q . It follows that we can fix an underlying smooth surface \mathbb{S} , and view the points of Q as defining pairs consisting of a complex structure on \mathbb{S} together with a meromorphic quadratic differential ϕ on the resulting Riemann surface S . Composing with a smoothly varying family of diffeomorphisms we can further assume that the differentials in Q have poles and zeroes at the same fixed points of \mathbb{S} .

Lemma 4.3. *Fix a constant $R > 1$. Then any point of Q is contained in some neighbourhood $U \subset Q$ such that*

$$(1/R) \cdot \ell_{\phi_1}(\gamma) \leq \ell_{\phi_2}(\gamma) \leq R \cdot \ell_{\phi_1}(\gamma),$$

for any curve γ in \mathbb{S} , and any pair of differentials $\phi_i \in U$.

Proof. Fix an arbitrary Riemannian metric g on the smooth surface \mathbb{S} , and write $\eta(x, y)$ for the distance between two points $x, y \in \mathbb{S}$ computed in this metric. Away from the poles p_i we can view the meromorphic differential ϕ corresponding to a point of Q as defining a smooth section of the bundle $(T_{\mathbb{S}}^*)^{\otimes 2}$, the tensor square of the rank 2 bundle of smooth complex-valued 1-forms on \mathbb{S} . Near a pole p_i of order m_i , the rescaled section $\eta(x, p_i)^{m_i} \cdot \phi(x)$ is smooth in a neighbourhood of p_i , and has non-zero value at p . Similar remarks apply near a zero of ϕ .

Given two points $\phi_1, \phi_2 \in Q$ it follows that the ratio $|\phi_1|/|\phi_2|$, considered as a smooth function on the set of nonzero tangent vectors to \mathbb{S} , is everywhere defined and varies

³More precisely, if ϕ_0 is an orbifold point, we should take an étale map $Q \rightarrow \text{Quad}(g, m)$ from a complex ball, but we suppress this point in what follows. Alternatively one could pull back the bundle (2.7) to Teichmüller space and work locally there.

smoothly with the differentials ϕ_i . Thus around any point of Q we can find a neighbourhood $U \subset Q$ such that $(1/R) \cdot |\phi_2| \leq |\phi_1| \leq R \cdot |\phi_2|$, for all $\phi_1, \phi_2 \in U$ and all tangent vectors to \mathbb{S} . Integrating this inequality along a curve gives the result. \square

4.3. Persistence of saddle connections. In this section we show that if a GMN differential varies continuously in its moduli space then its geodesics also vary continuously. We take notation as in the last section.

Proposition 4.4. *Suppose that $\gamma_0 \in \mathcal{C}(p, q)$ is a ϕ_0 -geodesic. Then there is a family of curves $\gamma(\phi) \in \mathcal{C}(p, q)$, varying continuously with $\phi \in Q$, such that $\gamma_0 = \gamma(\phi_0)$, and such that for all $\phi \in Q$ the curve $\gamma(\phi)$ is a ϕ -geodesic.*

Proof. Let us first consider the case when ϕ_0 has no simple poles. By Theorem 4.1, for each $\phi \in Q$ there exists a unique ϕ -geodesic $\gamma(\phi)$ in $\mathcal{C}(p, q)$ which is homotopic to γ_0 . We must show that the resulting curves $\gamma(\phi)$ vary continuously with ϕ . Assuming the opposite, let us take $\epsilon > 0$ and suppose that there exists a sequence of differentials $\phi_n \in Q$ with $\phi_n \rightarrow \phi$, such that for all n the geodesic $\gamma_n = \gamma(\phi_n)$ does not lie within distance ϵ of $\gamma = \gamma(\phi)$ in the supremum norm. In other words, for each n , we can find $t_n \in [0, 1]$ such that

$$d(\gamma_n(t_n), \gamma(t_n)) \geq \epsilon.$$

Passing to a subsequence we can assume that $t_n \rightarrow t \in [0, 1]$. Lemma 4.3 shows that for any $R > 1$

$$(4.1) \quad (1/R) \cdot \ell_\phi(\gamma_n) \leq \ell_{\phi_n}(\gamma_n) \leq \ell_{\phi_n}(\gamma) \leq R \cdot \ell_\phi(\gamma),$$

for large enough n . In particular, we can assume that the γ_n all satisfy $\ell_{\phi_n}(\gamma_n) \leq L$, for some constant $L > 0$. Theorem 4.1 implies that, when parameterised proportional to ϕ -arclength, some subsequence of the γ_n converges to a limit curve γ_L . This limit curve cannot be equal to γ , since

$$d(\gamma_L(t), \gamma(t)) \geq \epsilon.$$

On the other hand, the inequalities (4.1) show that $\ell_\phi(\gamma_L) \leq \ell_\phi(\gamma)$. This contradicts the fact, immediate from Theorem 4.1, that all geodesics are minimal.

For the general case we use the same trick as in Lemma 4.2. Namely, we consider a covering $\pi: \tilde{S} \rightarrow S$ which is branched at all simple poles of ϕ_0 . We can lift γ_0 to a geodesic $\tilde{\gamma}_0$ on the surface \tilde{S} for the pulled-back differential $\pi^*(\phi_0)$. This differential has no simple poles, so we can apply what we proved above to obtain a continuous deformation of $\tilde{\gamma}_0$. Pushing back down to S gives the required deformation of γ_0 . \square

Remarks 4.5. (a) *If the geodesic $\gamma_0 = \gamma(\phi_0)$ of Proposition 4.4 is a straight arc (which is to say that it contains no zeroes of ϕ_0 in its interior) then, by continuity, the same is true for the geodesics $\gamma(\phi)$ for all differentials ϕ in some neighbourhood of ϕ_0 . Thus, in particular, saddle connections persist under small deformations of the differential.*

- (b) *A minor modification of the proof shows that the conclusion of Proposition 4.4 also holds if we allow the endpoints p, q of the path $\gamma(\phi)$ to vary continuously with the differential ϕ .*

4.4. Persistence of separating trajectories. We explained in Section 3.3 that an infinite critical point p of a meromorphic quadratic differential is contained in a trapping neighbourhood $p \in U$ such that all trajectories entering U eventually tend towards the point p . In fact we can be more explicit about this neighbourhood.

Lemma 4.6. *Take a point $p \in \text{Crit}_\infty(\phi)$ which is not a double pole with real residue. Then there is a disc $p \in D \subset S$ whose boundary consists of saddle connections and such that any trajectory intersecting D tends to p in at least one direction.*

Proof. Consider the geodesic representative of the closed loop δ_p around p . It consists of a union of straight arcs of varying phase connecting zeroes of ϕ_0 , which together cut out an open disc $p \in D \subset S$ containing no points of $\text{Crit}_\infty(\phi)$. This disc cannot contain any finite critical points of ϕ either: if $z \in D$ were such a point, the geodesic representative of a loop round p based at z would be homotopic to δ_p , contradicting uniqueness of geodesic representatives. If a trajectory intersects the boundary of D twice this again contradicts uniqueness of geodesics. Hence any trajectory in one direction must either be recurrent or tend to the pole. But recurrence is also impossible since the boundary of the resulting spiral domain would involve saddle connections contained in D . \square

Remarks 4.7. (a) *In the case of a double pole with real residue, the pole is enclosed in a degenerate ring domain whose boundary consists of a union of saddle trajectories. This ring domain is the analogue of the trapping neighbourhood: any trajectory intersecting D is one of the closed trajectories of D .*

- (b) *Proposition 4.4 shows that the region $D = D(\phi)$ of Lemma 4.6 varies continuously with ϕ . In particular, there is an open neighbourhood of the pole p which has the trapping property for all differentials in a neighbourhood of a given base-point ϕ_0 .*

In the last section we proved that saddle connections persist to nearby differentials; we shall now prove a similar result for separating trajectories. Note that in contrast to saddle connections (whose phases vary as they deform) we can always deform separating trajectories in such a way that they remain horizontal.

Proposition 4.8. *Suppose that $\gamma_0: [0, \infty) \rightarrow S^\circ$ is a separating trajectory for the differential ϕ_0 , which starts at a point $p \in S^\circ$ and limits to an infinite critical point $r \in S$. Then there is a neighbourhood $\phi_0 \in U \subset Q$, and a family of curves $\gamma(\phi): [0, \infty) \rightarrow S^\circ$, varying continuously with $\phi \in Q$, such that $\gamma_0 = \gamma(\phi_0)$, and such that for all $\phi \in Q$ the curve $\gamma(\phi)$ is a separating trajectory for ϕ , starting at p and limiting to r .*

Proof. Note that r cannot be a double pole of real residue. Consider the open neighbourhood $r \in D \subset S$ which has the trapping property for any ϕ lying in some neighbourhood $\phi_0 \in U \subset Q$. Take a point $q_0 = q(\phi_0) \in D$ on the trajectory γ_0 . Consider the holomorphic function near q_0 obtained by integrating $\sqrt{\phi}$ along the trajectory γ_0 . This function varies smoothly with ϕ so, by the implicit function theorem, we can continuously vary $q(\phi) \in D$ so that

$$(4.2) \quad \int_p^{q(\phi)} \sqrt{\phi} \in \mathbf{R},$$

for all $\phi \in U$, where the integral is taken along a path homotopic to γ_0 .

By Remark 4.5(b), there is a continuous family of curves $\gamma(\phi)$ parameterized by $\phi \in U$, with $\gamma(\phi_0) = \gamma_0$, and such that for each ϕ the curve $\gamma(\phi)$ is a ϕ -geodesic connecting p to $q(\phi)$. Shrinking U if necessary, each of these geodesics is in fact a straight arc, and the relation (4.2) shows that these arcs are all horizontal. By the trapping assumption on D , each arc $\gamma(\phi)$ must extend to a separating trajectory $\gamma(\phi): [0, \infty) \rightarrow S$ for ϕ . The fact that these trajectories vary continuously when restricted to any finite interval $[0, t] \subset [0, \infty)$ then follows by another application of the argument of Proposition 4.4. \square

4.5. Horizontal strip decompositions. Fix again a genus $g \geq 0$ and a collection of $d \geq 1$ unordered positive integers $m = \{m_i\}$. As preparation for proving that the period map (2.10) is a local isomorphism, in this section and the next we will study the set of all saddle-free GMN differentials whose separating trajectories decompose the underlying smooth surface \mathbb{S} into a given fixed set of horizontal strips and half-planes.

We say that two saddle-free GMN differentials (S_i, ϕ_i) have *the same horizontal strip decomposition* if there is an orientation-preserving diffeomorphism $f: S_1 \rightarrow S_2$ which maps each horizontal strip (respectively half-plane) of ϕ_1 bijectively onto a horizontal strip (respectively half-plane) of ϕ_2 . In particular, equivalent differentials have the same horizontal strip decomposition.

More concretely, two equivalence-classes of saddle-free differentials have the same horizontal strip decomposition precisely if we can find representatives (S_i, ϕ_i) which have the same underlying smooth surface \mathbb{S} , and the same horizontal strips, half-planes and separating trajectories.

We would like to classify equivalence classes of saddle-free differentials (S, ϕ) with a given horizontal strip decomposition in terms of the periods of the corresponding standard saddle classes α_h . However, the existence of differentials with automorphisms which permute their horizontal strips makes it impossible to assign a well-defined period point to an arbitrary saddle-free differential. The solution is to consider framed differentials, as in Section 2.6.

We say that two framed GMN differentials have the same horizontal strip decomposition if there is an orientation-preserving diffeomorphism $f: S_1 \rightarrow S_2$ preserving the horizontal strip decomposition as before, and also preserving the framings, in the sense that the distinguished lift \hat{f} of Section 2.6 makes the diagram (2.8) commute. Again, equivalent framed differentials have the same horizontal strip decomposition.

Note that, by Lemma 3.2, a framing of a saddle-free differential gives rise to a labelling of the horizontal strips by the elements of a basis of Γ , and that conversely, the framing is completely determined by this labelling. Explicitly, if the framing is given by an isomorphism $\theta: \Gamma \rightarrow \hat{H}(\phi)$, then the strip h is naturally labelled by the element $\theta^{-1}(\alpha_h)$. Moreover, two saddle-free differentials have the same horizontal strip decomposition precisely if we can find representatives (S_i, ϕ_i) which have the same underlying smooth surface \mathbb{S} , and the same horizontal strips as before, and which moreover have the same labellings by elements of Γ .

The following result will be the basis for our proof of the existence of period co-ordinates. We defer the proof to the next subsection: by what was said above it amounts to classifying saddle-free differentials ϕ on a smooth surface \mathbb{S} with a fixed set of horizontal strips and half-planes, and also with a fixed ordering of the horizontal strips.

Proposition 4.9. *Let $U \subset \text{Quad}^\Gamma(g, m)$ be the set of equivalence-classes of framed saddle-free GMN differentials with a given horizontal strip decomposition. Choosing an ordering of the horizontal strips, the resulting map*

$$\pi_U: U \rightarrow \mathbb{C}^n, \quad \phi \mapsto Z_\phi(\alpha_{h_i})$$


is a bijection onto the subset $\{(z_1, \dots, z_n) \in \mathbb{C}^n : \text{Im}(z_i) > 0\}$.

The next example shows that it is possible for a saddle-free differential to have non-trivial automorphisms which preserve each horizontal strip. Such automorphisms preserve the standard arc classes and hence give automorphisms of the corresponding framed differential.

Example 4.10. *Consider the case $g = 1$ and $m = (2)$: one of the exceptional cases of Lemma 2.6. The space $\text{Quad}(g, m)$ parameterizes pairs (S, ϕ) , where S is a Riemann surface of genus 1, and ϕ is a meromorphic differential with one double pole and two simple zeroes. Such differentials can be written explicitly as*

$$\phi(z) = (a\wp(z) + b) dz^{\otimes 2},$$

where $\wp(z)$ is the Weierstrass \wp -function corresponding to S . These functions are invariant under the inverse map $z \mapsto -z$.


 FIGURE 17. Horizontal strip decompositions in case $g = 1$, $m = (2)$.

The possible horizontal strip decompositions are shown in Figure 17. Note that the inverse map (which is a rotation by π on the diagram) preserves each of these decompositions, and acts via a non-trivial automorphism of each horizontal strip.

4.6. Gluing strips. In this section we prove Proposition 4.9.

Let $\mathfrak{h} \subset \mathbb{C}$ be the upper half-plane, and take $z \in \mathfrak{h}$. We define the standard complete horizontal strip of period z to be the region

$$\bar{C}(z) = \{0 \leq \text{Im}(t) \leq \text{Im}(z)\} \subset \mathbb{C},$$

with two marked points on its boundary at $\{0, z\}$. We equip the interior $C(z) \subset \bar{C}(z)$ with the quadratic differential $dt^{\otimes 2}$. Similarly, the standard complete half-plane $\bar{C}(\infty)$ is the region $\{\text{Im}(t) \geq 0\}$, equipped with the differential $dt^{\otimes 2}$ in its interior, and with a single marked point at 0.

For any two elements $w, z \in \mathfrak{h}$ there is a diffeomorphism

$$\theta_{w,z}: \bar{C}(w) \cong \bar{C}(z),$$

preserving the marked points on the boundary, and with the further property that in a neighbourhood of each of the two boundary components of $\bar{C}(w)$ it is given by a translation in \mathbb{C} . To be completely definite, we can define

$$\theta_{w,z}(t) = t + \eta(\text{Im}(t)/\text{Im}(w)) \cdot (z - w),$$

where $\eta: [0, 1] \rightarrow [0, 1]$ is some smooth function satisfying $\eta([0, \frac{1}{4}]) = 0$ and $\eta([\frac{3}{4}, 1]) = 1$.

When $z \in \mathfrak{h}$ there is a single non-trivial automorphism of $\bar{C}(z)$ preserving the differential and the marked points, namely $t \mapsto z - t$. We can ensure that the diffeomorphisms $\theta_{w,z}$ we have constructed commute with these non-trivial automorphisms by insuring that the function η satisfies $\eta(t) + \eta(1 - t) = 1$.

Let ϕ be a saddle-free GMN differential on a compact Riemann surface S such that $\text{Crit}_\infty(\phi)$ is non-empty. Thus ϕ determines a decomposition of the underlying smooth surface \mathbb{S} into horizontal strips and half-planes. The restriction of the differential ϕ to a horizontal strip h_i is equivalent to the standard differential $dt^{\otimes 2}$ on the standard cell $C(z_i)$ via an isomorphism $f_i: C(z_i) \rightarrow h_i$. This extends to a continuous map

$$\bar{f}_i: \bar{C}(z_i) \rightarrow \mathbb{S},$$

and composing with a translation we can ensure that it takes the marked points $\{0, z_i\}$ to the finite critical points on the boundary of h_i . The four boundary half-edges of $\bar{C}(z_i)$ are then taken to the separating trajectories forming the boundary of h_i .

To build a differential ϕ on \mathbb{S} with the same horizontal strip decomposition, and arbitrary periods w_i , introduce diffeomorphisms

$$g_i = f_i \circ \theta_{w_i, z_i}: C(w_i) \rightarrow h_i.$$

Pushing forward the complex structure and quadratic differential from $C(w_i)$ using g_i defines a new complex structure and differential ψ on the strips h_i , and this trivially extends over the separating trajectories and finite critical points since it agrees with the old one ϕ in a neighbourhood of these points. Note that we leave the half-planes completely unchanged.

We must now show that the new complex structure extends over the poles of ϕ . First note that the function θ_{w_i, z_i} is invariant under translations in the real direction in $\bar{C}(w_i)$, and hence its derivatives are bounded on $\bar{C}(w_i)$. It follows that there is a bound

$$(4.3) \quad (1/R) \cdot |\phi| \leq |\psi| \leq R \cdot |\phi|,$$

where we consider both sides as functions on the tangent bundle to $S^\circ = S \setminus \text{Crit}_\infty(\phi)$, and $R > 1$ is a constant depending only on the periods w_i and z_i .

Take a small punctured disc $U \subset \mathbb{S}$ centered at a pole p of ϕ , and consider the complex structures $U(\phi)$ and $U(\psi)$ induced by the two differentials. Thus $U(\phi)$ is biholomorphic to the standard punctured disc D^* , and we would like to know that this is also the case for $U(\psi)$. By the Riemann mapping theorem, $U(\psi)$ is biholomorphic to some annulus

$$\{r_1 < |z| < r_2\} \subset \mathbb{C}.$$

We can compute the modulus $(1/2\pi) \log(r_2/r_1) \in [0, \infty]$ using extremal length [9], and the inequalities (4.3) show that this gives the same result as for $U(\phi)$. Hence $U(\psi)$ is also biholomorphic to a punctured disc, and so we can extend the new complex structure over p . Applying the inequalities (4.3) again then shows that ψ extends to a meromorphic function at p with the same pole order as ϕ .

The above argument proves that the map π_U of Proposition 4.9 is surjective. To prove that it is injective, suppose that two differentials (S_i, ϕ_i) have the same horizontal strip decomposition and the same periods $Z_{\phi_i}(\alpha_{h_j})$. Note that the restrictions of (S_i, ϕ_i) to the interior of a given horizontal strip h_j are equivalent, via a biholomorphism which extends continuously over the boundary of the strip. Glueing these maps together gives a homeomorphism $f: S_1 \rightarrow S_2$ which is biholomorphic on the interior of each strip. It follows that f is in fact a biholomorphism, and since the meromorphic sections $f^*(\phi_2)$ and ϕ_1 coincide on an open subset, they must be equal.

4.7. Period co-ordinates. We can now prove that (with certain exceptions) the period map (2.10) is a local isomorphism. Let us fix a genus $g \geq 0$ and a collection of $d \geq 1$ integers $m = \{m_i\}$. We shall need the following easy corollary of Proposition 4.8.

Lemma 4.11. *Suppose that at least one $m_i \geq 2$. Then the subset $B_0 \subset \text{Quad}(g, m)$ of saddle-free differentials is open and has non-trivial intersection with every S^1 -orbit.*

Proof. By Lemma 3.1, a GMN differential with an infinite critical point is saddle-free precisely if every trajectory leaving a finite critical point is separating. Proposition 4.8 shows that this condition is stable under small deformations of the differential. Thus B_0 is open.

If a GMN differential ϕ has a saddle trajectory γ then by the definition of the hat-homology class, $Z_\phi(\hat{\gamma}) \in \mathbf{R}_{>0}$. Consider the subset $\Theta_\phi \subset S^1$ of phases θ for which $e^{-i\pi\theta} \cdot \phi$ has a saddle trajectory. Then Θ_ϕ is contained in the set of elements $\arg Z_\phi(\alpha)$ for classes $\alpha \in \hat{H}(\phi)$ having nonzero period. Thus Θ_ϕ is countable. In particular, the complement of Θ_ϕ is non-empty. \square

We remark that the conclusion of Lemma 4.11 is definitely false without the assumption that the differential has an infinite critical point: in the case of finite-area differentials on a compact surface, Masur [27] proved that the (still countable) subset $\Theta_\phi \subset S^1$ is dense.

Theorem 4.12. *Suppose that the polar type (g, m) is not one of the 6 exceptional cases listed in Lemma 2.6. Then the space of framed GMN differentials $\text{Quad}^\Gamma(g, m)$ is either empty, or is a complex manifold of dimension n , and the period map*

$$\pi: \text{Quad}^\Gamma(g, m) \rightarrow \text{Hom}_{\mathbb{Z}}(\Gamma, \mathbb{C})$$

is a local isomorphism of complex manifolds.

Proof. We divide into two cases. In the finite area case when all $m_i = 1$ we appeal to the known result that the period map is a local isomorphism in this setting.⁴ We do need to be a little bit careful to prove that $\text{Quad}^\Gamma(g, m)$ is a manifold. Consider the complex manifold $\mathcal{V}(g, m)$ obtained by pulling back the fibration

$$\text{Quad}(g, m) \rightarrow \mathcal{M}(g, d) / \text{Sym}(m)$$

to Teichmüller space. Pulling back the space of framed differentials gives a local isomorphism $p: \mathcal{V}^\Gamma(g, m) \rightarrow \mathcal{V}(g, m)$. Taking an open subset $U \subset \mathcal{V}(g, m)$ and a locally-defined section of p gives local isomorphisms

$$f: U \rightarrow \mathcal{V}^\Gamma(g, m) \rightarrow \text{Quad}^\Gamma(g, m),$$

and taking the composition with the period map π gives a locally-defined period map on $\mathcal{V}(g, m)$. This period map is known to be a local isomorphism [36] and, hence, shrinking

⁴In fact this appeal can be avoided: see Remark 6.5.

U if necessary, we can assume that f is injective on points. Our assumption ensures that the space $\text{Quad}^\Gamma(g, m)$ has trivial generic automorphism group so it follows that f is an isomorphism onto its image. Hence $\text{Quad}^\Gamma(g, m)$ is a complex manifold and the period map π is a local isomorphism.

Suppose now that some $m_i \geq 2$. By Lemma 4.11 we can use the S^1 -action and work in a neighbourhood consisting of saddle-free differentials on a fixed underlying surface \mathbb{S} and with a fixed horizontal strip decomposition. An automorphism of such a differential ϕ is a smooth map $f: \mathbb{S} \rightarrow \mathbb{S}$ satisfying $f^*(\phi) = \phi$. It preserves the framing precisely if it acts trivially on the set of horizontal strips. Assume that f is not the identity. When pulled-back to a standard strip $C(z_i)$ it must then act by $t \mapsto z_i - t$. But the construction of Section 4.6 shows that f then preserves all differentials with the same horizontal strip decomposition as ϕ . Since $\text{Quad}^\Gamma(g, m)$ is assumed to have trivial generic automorphism group, this is impossible. Hence $\text{Quad}^\Gamma(g, m)$ is a manifold. The result now follows from Proposition 4.9. The horizontal strip decomposition is locally-constant on B_0 , so the subset U appearing there is open. The map π_U is certainly holomorphic because its components are periods of the spectral cover, which varies holomorphically with ϕ . Since π_U is also bijective, it is an isomorphism. \square

5. STRATIFICATION BY NUMBER OF SEPARATING TRAJECTORIES

This rather technical section contains some further results concerning the trajectory structure of GMN differentials. We focus particularly on the stratification of the space $\text{Quad}(g, m)$ by differentials with a fixed number of separating trajectories. Throughout, we fix a genus $g \geq 0$ and a polar type $m = \{m_i\}$ such that all $m_i \geq 2$, and consider differentials in $\text{Quad}(g, m)$. In particular, all differentials are complete and have at least one infinite critical point.

5.1. Homology classes of saddle trajectories. Let ϕ be a complete GMN differential on a Riemann surface S . Recall from Section 3.1 that every saddle trajectory γ has an associated hat-homology class $\hat{\gamma} \in \hat{H}(\phi)$. We say that two saddle trajectories γ_1 and γ_2 are *hat-homologous* if $\hat{\gamma}_1 = \hat{\gamma}_2$. More generally, we say that γ_1 and γ_2 are *hat-proportional* if

$$\mathbb{Z}_{>0} \cdot \hat{\gamma}_1 = \mathbb{Z}_{>0} \cdot \hat{\gamma}_2 \subset \hat{H}(\phi).$$

The following result, which is the analogue in our situation of a result of [28], relies essentially on the assumption that all finite critical points of ϕ are simple zeroes.

Lemma 5.1. *Suppose that (S, ϕ) admits a pair of distinct hat-proportional saddle trajectories γ_1, γ_2 . Then one of the following cases holds:*

- (i) *The γ_i are hat-homologous and closed and form the two boundary components of a non-degenerate ring domain.*

- (ii) *The surface S contains a separating non-degenerate ring domain A , bounding on one side an open genus one subsurface $T \subset S$ containing no critical points of ϕ . The boundary component of A adjoining T is a union of two non-closed saddle trajectories ν_1, ν_2 and either*
 - (iia) *the second boundary of A consists of more than one saddle trajectory; then $\{\gamma_1, \gamma_2\} = \{\nu_1, \nu_2\}$ and the γ_i are hat-homologous;*
 - (iib) *the second boundary of A is a closed saddle trajectory μ ; then $2\hat{\nu}_1 = \hat{\mu} = 2\hat{\nu}_2$ and $\{\gamma_1, \gamma_2\} \subset \{\mu, \nu_1, \nu_2\}$.*
- (iii) *The surface S is a torus and ϕ has a unique pole p , which has order 2 and real residue. The γ_i are hat-homologous, have distinct endpoints, and together form the boundary of the degenerate ring domain enclosing p .*

The three cases are illustrated in Figure 18. In the cases (ii) and (iii) the torus subsurface carries an irrational foliation, i.e. the interior of the torus is a spiral domain. The final paragraph of Section 6.4 explains a construction of quadratic differentials illustrating these cases.




FIGURE 18. Configurations of hat-proportional saddle trajectories.

Proof. We divide the proof into three cases, depending on how many of the saddle trajectories γ_i are closed.

Case (1). First suppose that γ_1 and γ_2 are both closed. They must then be disjoint. Each γ_i is one boundary component of a ring domain A_i . If both A_i are degenerate, then $\hat{\gamma}_i = \beta_{p_i}$ for different double poles p_i , and these classes are linearly independent by Lemma 2.3. Thus we can assume that A_1 is non-degenerate. Then, as in Figure 19, we can write $\hat{\gamma}_1 = \hat{\alpha}_+ - \hat{\alpha}_-$, with α_{\pm} being saddle trajectories contained entirely in the ring domain A_1 , and satisfying $\hat{\alpha}_+ \cdot \hat{\alpha}_- = \pm 2$. But then $\hat{\alpha}_+ \cdot \hat{\gamma}_1 = \pm 2$ also, so if the γ_i are hat-proportional, then γ_2 must meet α_+ , and hence must be the other boundary

component of A_1 . It then follows that the γ_i are hat-homologous, and this is case (i) of the Lemma.




FIGURE 19. A non-degenerate ring domain with a pair of arcs α_{\pm} .

Case (2). Next suppose that neither γ_1 nor γ_2 is closed. If $\gamma_1 \cup \gamma_2$ does not separate the surface S then we can find a path α in S connecting poles of ϕ , whose interior lies in $S \setminus \text{Crit}(\phi)$, and which intersects γ_1 once, and γ_2 not at all. If we take $\hat{\alpha}$ to be one of the two inverse images of this path on the spectral cover, then the Lefschetz pairings are $\langle \hat{\gamma}_1, \hat{\alpha} \rangle = \pm 1$ and $\langle \hat{\gamma}_2, \hat{\alpha} \rangle = 0$. Hence the γ_i are not hat-proportional.

Suppose then that γ_1 and γ_2 both have the same pair of endpoints $z_1 \neq z_2$, and together form a separating loop γ . Consider the third trajectories coming out of z_1 and z_2 . If these lie on opposite sides of γ then it is easily seen that $\hat{\gamma}_1 \cdot \hat{\gamma}_2 = \pm 2$ so again the γ_i are not hat-proportional. Thus we conclude that the loop γ is one boundary component of a ring domain A .

If there are poles on both sides of γ one can take a path between such poles meeting γ_1 exactly once and disjoint from γ_2 , showing the γ_i are not hat-proportional. Therefore we may assume that γ separates S , and bounds on one side a subsurface T containing no poles of ϕ . If T contains a zero q of ϕ , we may take a path α from a pole to itself which crosses γ_1 , encircles q once, and returns parallel to itself crossing γ_1 again, with α globally disjoint from γ_2 . Let $\hat{\alpha}$ denote one component of the preimage of α on the spectral cover. Then $\langle \hat{\alpha}, \hat{\gamma}_1 \rangle = \pm 2$ and $\langle \hat{\alpha}, \hat{\gamma}_2 \rangle = 0$, so the γ_i are not hat-proportional.

Since the other boundary component of the ring domain A contains zeroes or a double pole of ϕ , it necessarily lies outside T . Thus a closed trajectory σ inside A bounds a subsurface $R \supset T$ which contains two zeroes (lying on γ) and no poles. Doubling R along α then gives a holomorphic quadratic differential on a closed surface with exactly four simple zeroes, which implies that the subsurface R is a torus.

If the ring domain A is non-degenerate we are now in the setting of case (ii) of the Lemma, and the classes $\hat{\gamma}_i$ coincide since the preimage of γ bounds an unpunctured subsurface of \hat{S} . On the other hand, if A is a degenerate ring domain centered on a double pole p then we are in the setting of case (iii), and the γ_i are hat-homologous with hat-homology class equal to half the residue class β_p .

Case (3). Reordering the γ_i if necessary, we may now suppose that γ_1 is closed, and γ_2 has distinct endpoints. If γ_1 is not separating, there is a path from a pole to a pole which meets γ_1 once and is disjoint from γ_2 , so the γ_i are not hat-proportional. The same argument applies if γ_1 separates S into subsurfaces each of which contain poles, so we may assume that γ_1 bounds a subsurface R containing no pole of ϕ .

If the interior of γ_2 lies inside $S \setminus R$ then we may take a path α which goes from a pole, around one end-point of γ_2 , and back to the pole parallel to itself, and which is entirely disjoint from γ_1 . The Lefschetz pairing argument as above then implies that the γ_i are not hat-proportional. We may therefore assume that γ_2 lies inside the subsurface R bound by γ_1 .

If R contains some zero q not lying on either γ_i , we may pick a closed path α , disjoint from γ_2 , which starts at a pole of ϕ , crosses γ_1 , encircles q , and then returns parallel to itself crossing γ_1 . The Lefschetz pairings again imply that the γ_i are not hat-proportional. We may therefore assume that R contains no zeroes other than those lying on $\gamma_1 \cup \gamma_2$.

The curve γ_1 is a boundary component of some ring domain A . Suppose A is non-degenerate, and A is contained inside R . Then γ_2 is disjoint from γ_1 . Let σ be a closed trajectory in A . Doubling along σ yields a surface containing no poles and four simple zeroes, hence of genus 2, which implies that γ_1 bounds a torus, and γ_2 connects the two zeroes z_1, z_2 lying inside that torus. If both zeroes z_i lie on the boundary of A then, by considering intersections with closed curves in the torus, it is easy to see that γ_2 must be contained in the boundary of A . We are then in case (iib) of the Lemma. Otherwise, the other boundary component of A is comprised of a single saddle ν , and γ_2 meets ν in a point, giving a situation as on the left of Figure 20.




FIGURE 20. Configurations of closed curves which cannot arise as saddle trajectories.

We claim this configuration cannot occur. Indeed, if it could, one could replace everything outside σ with a degenerate ring domain, yielding the right-hand picture of Figure 20. But, as in Example 12.4, any differential on a torus with a single double pole

and simple zeroes is invariant under an involution of the torus which permutes the two zeroes, and this rules out the asymmetric trajectory structure shown.⁵

Suppose next that the ring domain A is non-degenerate, but not contained inside R . The third half-edge at the zero on γ_1 then enters R , and we can double along a closed trajectory σ in A to get a surface with at most six zeroes and no poles, hence with at most four zeroes. It follows again that R is a torus, and that γ_2 must intersect γ_1 at one point. This gives the same local configuration of saddles as in the previous case, and by gluing in a degenerate ring domain along σ one obtains the same contradiction as before.

Finally, if A is degenerate, one arrives directly at the second picture of Figure 20, and that again yields a contradiction. This then completes the proof. \square

5.2. Stratification. Let ϕ be a GMN differential on a Riemann surface S defining a point in $\text{Quad}(g, m)$. Note that we are assuming that all $m_i \geq 2$ so ϕ has no simple poles, and at least one infinite critical point. Since exactly 3 horizontal trajectories emerge from each zero of ϕ , there is an equality

$$r_\phi + 2s_\phi + t_\phi = k, \quad k = 3 |\text{Zer}(\phi)| = 3(4g - 4 + \sum_{i=1}^d m_i),$$

where r_ϕ is the number of trajectories that are recurrent in one direction, but tend to a zero in the other, s_ϕ is the number of saddle trajectories, and t_ϕ is the number of separating trajectories. Define subsets

$$B_p = \{\phi \in \text{Quad}(g, m) : r_\phi + 2s_\phi \leq p\}$$

Note that $B_0 = B_1$ is precisely the set of saddle-free differentials. Indeed, by Lemma 3.1, a differential having no saddle trajectories has no recurrent trajectories either. We call the elements of B_2 *tame differentials*; such differentials have at most one saddle trajectory.

Lemma 5.2. *The subsets $B_p \subset \text{Quad}(g, m)$ form an increasing chain of dense open subsets*

$$B_0 = B_1 \subset B_2 \subset \cdots \subset B_k = \text{Quad}(g, m).$$

Proof. This is very similar to the proof of Lemma 4.11. Since B_p is the subset of differentials for which $t_\phi \geq k - p$, the statement that B_p is open is equivalent to the condition that the function t_ϕ is lower semi-continuous. This follows from Proposition 4.8. If a differential ϕ has a saddle trajectory γ then $Z_\phi(\hat{\gamma}) \in \mathbf{R}$ by the definition of

⁵Alternatively, one could collapse a zero into the double pole using the local surgery from Section 6.4 below. This would yield a quadratic differential on a torus with a simple pole and a single zero, but as in Example 2.5, no such exists.

the hat-homology class. In local period co-ordinates the complement of B_0 is therefore contained in a countable union of real hyperplanes. \square

Define $F_p = B_p \setminus B_{p-1}$ for $p \geq 1$, and set $F_0 = B_0$. There is a finite stratification

$$\text{Quad}(g, m) = \bigsqcup_{p=0}^k F_p$$

by the locally-closed subsets F_i . The stratum F_1 is empty, and differentials in F_2 have exactly one saddle trajectory.

We call a GMN differential ϕ *generic* if the periods of non-proportional elements of the lattice $\hat{H}(\phi)$ define distinct rays in \mathbb{C} . More precisely, the condition is that for all $\gamma_1, \gamma_2 \in \hat{H}(\phi)$ there is an implication

$$\mathbf{R} \cdot Z_\phi(\gamma_1) = \mathbf{R} \cdot Z_\phi(\gamma_2) \implies \mathbb{Z}\gamma_1 = \mathbb{Z}\gamma_2.$$

It is easy to see that generic differentials are dense in $\text{Quad}(g, m)$: in local period co-ordinates the complement of the set of such differentials is contained in a countable union of real submanifolds cut out by relations of the form $Z_\phi(\gamma_1)/Z_\phi(\gamma_2) \in \mathbf{R}$.

We say that a differential ϕ is *0-generic* if the sublattice

$$\{\gamma \in \hat{H}(\phi) : Z_\phi(\gamma) \in \mathbf{R}\} \subset \hat{H}(\phi)$$

has rank ≤ 1 . This implies in particular that all saddle trajectories for ϕ are hat-proportional. Clearly, a differential ϕ is generic precisely if all elements of its S^1 -orbit are 0-generic.

5.3. Perturbing saddle trajectories. Let ϕ_0 be a GMN differential on a Riemann surface S defining a point

$$\phi_0 \in F_p \subset \text{Quad}(g, m)$$

for some $p \geq 2$ (recall that F_1 is empty). Our aim in this section and the next is to show that in period co-ordinates in a neighbourhood of ϕ_0 , the closed subset $F_p \subset B_p$ is contained in a real hyperplane. We begin by considering the case when ϕ_0 has saddle trajectories lying in the boundary of a horizontal strip or half-plane.

Proposition 5.3. *Suppose that $\phi_0 \in F_p$ has a half-plane or horizontal strip with a boundary component containing precisely $s \geq 1$ saddle trajectories γ_i . Let*

$$\alpha = \sum_{i=1}^s \hat{\gamma}_i \in \hat{H}(\phi_0)$$

be the sum of the corresponding hat-homology classes. Then there is an open neighbourhood $\phi \in U \subset B_p$ such that

$$\phi \in U \cap F_p \implies Z_\phi(\alpha) \in \mathbf{R}.$$

Proof. Considered as a subset of S , the half-plane or horizontal strip h is an open disc whose boundary is a closed curve (not necessarily embedded) made up of saddle trajectories and separating trajectories of ϕ_0 . By Propositions 4.4 and 4.8, if U is small enough, these trajectories deform continuously with the differential $\phi \in U$. The resulting deformed curve therefore also cuts out a disc in the surface S .

Integrating $\sqrt{\phi}$ inside this region gives a conformal mapping into \mathbb{C} which is a continuous perturbation of the horizontal strip or half-plane h . The boundary of the image region in




FIGURE 21. Perturbing a horizontal strip.

\mathbb{C} consists of straight lines connecting the images of the critical points of the differential. If the region in question is a horizontal strip there are two boundary components; composing with the map $z \mapsto -z$ we may assume that the saddle trajectories γ_i occur in the lower one.

Order the saddle trajectories from left to right (i.e. in anti-clockwise order around the boundary) and define real numbers

$$y_i = \operatorname{Im} Z_\phi(\hat{\gamma}_1 + \cdots + \hat{\gamma}_i), \quad 1 \leq i \leq s.$$

These numbers give the height of the vertices of the boundary of the perturbed strip, relative to the first vertex. In particular $y_s = \operatorname{Im} Z_\phi(\alpha)$. Note that the class α is definitely non-zero since $Z_{\phi_0}(\hat{\gamma}_j) \in \mathbf{R}_{>0}$.

Suppose that $\phi \in U \cap F_p$. This implies that if a horizontal arc emerging from a zero forms part of a non-separating trajectory for ϕ_0 , then the same must be true for the corresponding arc in ϕ . Working from the left, the first vertex with positive height y_i has a ray escaping to the pole on the left, which previously formed a saddle trajectory (see Figure 21). Thus we must have $y_i \leq 0$ for all i . Given this, if we also have $y_s < 0$ then the last vertex with $y_i = 0$ has a ray escaping to the pole on the right; if none of the vertices has height $y_i = 0$ then the very first vertex has such a ray. We conclude that we must also have $y_s = 0$. \square

Note that we actually proved more, namely that if $\phi \in U \cap F_p$ then $y_j \leq 0$ for $1 \leq j \leq s$.

5.4. Saddle reduction. As in the last section, let ϕ_0 be a GMN differential on a Riemann surface S defining a point

$$\phi_0 \in F_p \subset \operatorname{Quad}(g, m)$$

for some $p \geq 2$. We shall call a saddle-connection *borderline* if it lies in the boundary of a horizontal strip, half-plane or degenerate ring-domain.

The next result is analogous to Proposition 5.3 and deals with the case of saddle trajectories lying in the boundary of a degenerate ring-domain.

Lemma 5.4. *Suppose that $\phi_0 \in F_p$ contains a degenerate ring domain A centered on a double pole p . Then there is an open neighbourhood $\phi_0 \in U \subset B_p$ such that*

$$\phi \in U \cap F_p \implies Z_\phi(\beta_p) \in \mathbf{R}.$$

Proof. We can choose U so that we can reach any point by first deforming ϕ_0 maintaining the condition $Z_\phi(\beta_p) \in \mathbf{R}$, and then applying the S^1 -action. When $Z_\phi(\beta_p) \in \mathbf{R}$ the pole p still lies in a degenerate ring domain. Thus it is enough to deal with rotations. The boundary of A consists of a union of saddle trajectories. To understand trajectories for the rotated differential it is equivalent to consider non-horizontal trajectories for ϕ . It is clear that some of these will fall into the pole p . \square

Consider the closed subsurface with boundary $S_+ \subset S$ which is the closure of the union of the horizontal strips, half-planes and degenerate ring domains. Consider also the complementary closed subsurface $S_- \subset S$ which is the closure of the union of the spiral domains and non-degenerate ring domains. It is easy to see that these two surfaces S_\pm meet along a collection of simple closed curves made up of borderline saddle-connections.

Note that all infinite critical points of ϕ_0 are contained in the interior of S_+ , and since a GMN differential has a non-empty collection of poles, and we are assuming that all $m_i \geq 2$, it follows that S_+ is non-empty.

Proposition 5.5. *Take $p \geq 2$ and fix a point $\phi_0 \in F_p \subset \text{Quad}(g, m)$. Then there is a neighbourhood $\phi_0 \in U \subset B_p$ and a nonzero class $\alpha \in \hat{H}(\phi_0)$ such that*

$$\phi \in U \cap F_p \implies Z_\phi(\alpha) \in \mathbf{R}.$$

Proof. Since $p \geq 2$ there is at least one saddle trajectory for ϕ_0 . It follows that there must be at least one borderline saddle trajectory. Indeed, any saddle trajectory in S_+ is borderline, and if S_- is non-empty then S_- and S_+ are separated by borderline saddle trajectories. Combining Propositions 5.3 and 5.4 therefore gives the result. \square

It follows that, shrinking U if necessary, we can find a constant $r > 0$ such that

$$e^{i\pi\theta} \cdot \phi \in B_{p-1} \text{ when } 0 < |\theta| < r \text{ and } \phi \in U \cap F_p.$$

Thus we can always move to a larger stratum by small rotations of the differential.

5.5. Ring-shrinking. The assumption that a point $\phi \in \text{Quad}(g, m)$ is generic gives no restriction on which stratum F_p the differential ϕ lies in: although all saddle trajectories are hat-proportional, ϕ could well have a ring domain dividing the surface into two parts, one containing all the poles, and the other consisting of a spiral domain containing some large number of recurrent trajectories. For this reason it will be important in what follows to use the construction of Section 3.8 to eliminate ring-domains by shrinking them to a closed curve.

Recall that a ring domain is strongly non-degenerate if its boundary consists of two pairwise disjoint, simple, closed curves. The *width* of a non-degenerate ring domain is the minimal length of a path connecting the two boundary components. The width is a strictly positive real number; by a ring domain of width zero we mean any simple closed curve which is a union of saddle trajectories, and which is not a boundary component of a ring domain of strictly positive width. The following result, which will be used in Section 5, shows that any strongly non-degenerate ring domain may be shrunk to width zero. The result is illustrated in Figure 22.




FIGURE 22. Shrinking a ring domain to have width zero.

Proposition 5.6. *Suppose that a differential $(S_1, \phi_1) \in \text{Quad}(g, m)$ contains a strongly non-degenerate ring domain A of width $w > 0$. Then there is a continuous family $(S_t, \phi_t) \in \text{Quad}(g, m)$ parameterized by $t \in [0, 1]$, such that each surface S_t contains a ring domain A_t of width $t \cdot w$, strongly non-degenerate if $t > 0$, and there are equivalences*

$$(S_t \setminus \bar{A}_t, \phi_t|_{S_t \setminus \bar{A}_t}) \cong (S \setminus \bar{A}, \phi|_{S \setminus \bar{A}}).$$

Proof. The non-degenerate ring domain A is equivalent to a region

$$\{a < |z| < b\} \subset \mathbb{C} \text{ equipped with } \phi(z) = r \cdot dz^{\otimes 2}/z^2$$

for some $r \in \mathbf{R}_{<0}$. The only invariants are the width, which is $w = \log(b/a)$, and the length of the two boundary components, which is $2\pi\sqrt{r}$. For $t \in (0, 1)$ we define A_t to be the ring domain with the same length boundary components as A , but with width $t \cdot w$. We define the surface S_t by glueing A_t into $S \setminus A$ using Lemma 3.3. There is a choice of glueing, since one may rotate one boundary component relative to the other by an angle θ . The end-point surface S_0 is again constructed using Lemma 3.3, by directly glueing the two components of $S \setminus A$. To ensure that the resulting differential ϕ_0 has simple zeroes we may need to take the rotation parameter θ to be nonzero. \square

5.6. Walls have ends. We have shown above that for $p \geq 2$ the stratum $F_p \subset \text{Quad}(g, m)$ is contained in a real hyperplane in local period co-ordinates. It can therefore be thought of as a wall, potentially dividing two different connected components of the open subset B_{p-1} . Now we want to go one step further and show that if $p > 2$ then these walls always have ends: we can move along the stratum F_p to get to a point near which the subset B_{p-1} is locally connected.

Proposition 5.7. *Assume that the polar type is not $m = (2)$ and take $p > 2$. Then every connected component of $F_p \subset \text{Quad}(g, m)$ contains a point ϕ with a neighbourhood $\phi \in U \subset B_p$, as in Proposition 5.5, such that $U \cap B_{p-1}$ is connected.*

Proof. Take $\phi_0 \in F_p$ and a neighbourhood U as in Proposition 5.5. Consider the inclusion

$$U \cap F_p \subset \{\phi \in U : Z_\phi(\alpha) \in \mathbf{R}\}$$

If this inclusion is strict, the wall has a hole in it, and then $U \cap B_{p-1}$ is connected and we are done. Otherwise, these two subsets are equal, and so staying in the same connected component of F_p we can replace ϕ_0 with a very close generic differential ϕ_1 .

Suppose that ϕ_1 has only one saddle trajectory γ . Then, since $p > 2$, there must exist recurrent trajectories. The surfaces S_\pm introduced in Section 5.4 are thus both non-empty, and must meet along γ . Then γ is closed and forms one boundary component of a ring domain A , which has to be degenerate, since there are no saddle trajectories to form its other boundary. Thus we conclude that $S_+ = A$, and since all poles of ϕ_1 lie in S_+ it follows that ϕ_1 has a single pole p of order 2, and that α is the corresponding residue class β_p .

Suppose then that ϕ_1 has more than one saddle trajectory. By the genericity assumption these are all hat-proportional, so they are arranged as in one of the cases of Lemma 5.1. It follows that there are two possibilities, corresponding to cases (i) and (iia) of Lemma 5.1: case (iii) is ruled out by the assumption on the polar type, and case (iib) cannot occur for a 0-generic differential, since not all saddle trajectories appearing are hat-proportional. In particular we see that ϕ_1 has a unique ring domain A , which is strongly non-degenerate, and whose boundary consists of either 2 or 3 saddle trajectories.

By Proposition 5.6, we can move along a path in $\text{Quad}(g, m)$ in which A shrinks so as to have width 0, but the rest of the differential remains unchanged. It is clear that this path remains in the stratum F_p since any separating trajectory lies outside A and hence is unaffected by the shrinking process. At the end of this operation we arrive at a differential ϕ_2 with no closed trajectories and either 2 or 3 saddle trajectories γ_i , which together form a simple closed curve γ .

We claim that all the saddle trajectories γ_i are borderline. Indeed, if the surface S_- of Section 5.4 is empty then all saddle trajectories are borderline, and otherwise the two surfaces S_\pm are separated by a simple closed curve made up of saddle trajectories,

which must be γ . Examining the configuration of trajectories near γ in the two cases it is easy to see that exactly two of the γ_i must lie in the boundary of a single horizontal strip or half-plane h . Proposition 5.3, and the remark following it, shows that there is a neighbourhood $\phi_2 \in U \subset B_p$ such that, with appropriate ordering of the γ_i ,

$$\phi \in U \cap F_p \implies y = \operatorname{Im} Z_\phi(\hat{\gamma}_1 + \hat{\gamma}_2) = 0 \text{ and } z = \pm \operatorname{Im} Z_\phi(\hat{\gamma}_1) \leq 0.$$

Lemma 5.1 shows that the saddle trajectories γ_1, γ_2 are not hat-proportional, so the variables y and z form part of a co-ordinate system near ϕ_2 . It follows that $U \setminus F_p$ is locally connected near ϕ_2 . \square

5.7. Homotopies to tame paths. In the proof of our main Theorems we shall need the following consequence of Proposition 5.7.

Proposition 5.8. *Assume that the polar type is not $m = (2)$. Then any path β in $\operatorname{Quad}(g, m)$ connecting two points of B_2 is homotopic relative to its end-points to a path in B_2 .*

Proof. Let us inductively assume that β has been deformed so as to lie in B_p for some $p > 2$. By Proposition 5.5 we can cover β by open subsets in which $F_p \subset B_p$ is contained in a real hyperplane. We can then wiggle it a little so that it meets F_p at a finite number of points ϕ_i . We now show how to deform β so as to reduce the number k of these points. Repeating the argument, we can deform β to lie in B_{p-1} . The result then follows by induction.

To eliminate a point $\phi = \phi_i$ we first use Proposition 5.7 to construct a path δ in F_p connecting ϕ to a point ψ where B_{p-1} is locally connected. Consider paths δ_\pm obtained by small rotations of δ in opposite directions. By Proposition 5.5 these can be assumed to lie entirely in B_{p-1} . Inserting these paths into β we obtain a homotopic path which crosses F_p at the point ψ instead of ϕ . Since B_{p-1} is locally connected near ψ we can then deform β further and so eliminate one of its intersections with F_p . \square

The assumption on the polar type in Propositions 5.7 and 5.8 is necessary, as we explain in the following remark.

Remark 5.9. *Suppose that the polar type is $m = (2)$ and consider the holomorphic function*

$$Z_\phi(\beta_p)^2: \operatorname{Quad}(g, m) \rightarrow \mathbb{C}^*.$$

We claim that this maps the subset B_2 into the complement of $\mathbf{R}_{>0} \subset \mathbb{C}^$. Thus for paths in B_2 the function $Z_\phi(\beta_p)^2$ does not wind around the origin. But by rotating a differential, it is easy to construct paths in $\operatorname{Quad}(g, m)$ for which this function does wind around the origin. Thus it follows that Proposition 5.8, and hence also Proposition 5.7, are false in this case.*

To prove the claim note that if $\phi \in B_2$ and $Z_\phi(\beta_p) \in \mathbf{R}$ then the unique pole p is contained in a degenerate ring domain A . The boundary of A must then be a single closed saddle trajectory, and the third trajectory leaving the zero on the boundary cannot be a saddle trajectory, or a separating trajectory, or divergent. This gives a contradiction.

5.8. More on ring-shrinking. We assume in this section that if $g = 1$ then the polar type is not $m = (2)$. Suppose that $\phi_+ \in \text{Quad}(g, m)$ is a 0-generic differential with more than one saddle trajectory. As in the proof of Proposition 5.7, it follows that ϕ_+ has a unique ring domain A , which is moreover strongly non-degenerate, and we can shrink A to obtain a differential ϕ with a closed curve γ formed of a union of either 2 or 3 non-closed saddle trajectories γ_i .

Note that the γ_i are the only saddle trajectories for ϕ . Let us write $\alpha_i = \hat{\gamma}_i \in \hat{H}(\phi)$. Examining Figure 18, it is easily seen that we can order the γ_i so that the complete set of finite-length trajectories for ϕ_+ in the two cases is as follows:

- (J1) a single ring domain A of class $\alpha = \alpha_1 + \alpha_2$ whose boundary components are closed saddle trajectories of the same class;
- (J2) a single ring domain A of class $\alpha = \alpha_1 + \alpha_2 + \alpha_3$, one of whose boundary components is a closed saddle trajectory of the same class, the other being a union of two non-closed saddle trajectories of equal classes $\alpha_1 + \alpha_2$ and α_3 .

The labelling of the γ_i is completely determined if we insist that

$$(5.1) \quad \text{Im } Z_{\phi_+}(\alpha_1)/Z_{\phi_+}(\alpha_2) > 0,$$

and we shall always follow this convention. Note that in the case (J2) there is a relation $\alpha_1 + \alpha_2 = \alpha_3$.

Proposition 4.8 implies that for any differential sufficiently close to ϕ there are saddle connections deforming each of the saddle trajectories γ_i . We shall need the following statement later.

Proposition 5.10. *Given a class $\beta \in \hat{H}(\phi)$, there is a neighbourhood $\phi \in U \subset \text{Quad}(g, m)$ with the following property: if $\phi_- \in U$ satisfies*

$$(5.2) \quad \text{Im } Z_{\phi_-}(\alpha_1)/Z_{\phi_-}(\alpha_2) < 0,$$

and γ is a saddle trajectory for ϕ_- with hat-homology class β , then $\gamma = \gamma_i$, for some i , and hence $\beta = \alpha_i$.

Proof. Consider first the abstract situation in which two saddle trajectories γ_1, γ_2 for a differential ϕ meet at a zero z . Assume that the γ_i have non-proportional hat-homology classes α_i , and consider differentials on either side of the wall

$$\text{Im } Z_\phi(\alpha_1)/Z_\phi(\alpha_2) = 0.$$

As above there are saddle connections deforming each γ_i . Consider the union $\gamma_1 \cup \gamma_2$ near the zero z . Local calculations (see Figure 15) show that on one side of the wall this path is a geodesic, whereas on the other side it is not, since there is a shorter path which bypasses the zero z .

Consider now the differential ϕ obtained by shrinking the ring domain in ϕ_+ . The walls

$$\operatorname{Im} Z_\phi(\alpha_i)/Z_\phi(\alpha_j) = 0$$

for $i \neq j$ all coincide. On the side of this wall defined by (5.1), none of the unions $\gamma_i \cup \gamma_{i+1}$ is a geodesic, since the shortest paths in these homotopy classes cross the ring domain. It follows that on the side of the wall defined by (5.2) each of these unions is a geodesic.

Suppose for a contradiction that we can find a sequence of differentials ϕ_i satisfying (5.2), each with a saddle connection C_i of class β , and which tend to ϕ . The length of the saddle connections $\ell_{\phi_i}(C_i) = |Z_{\phi_i}(\beta)|$ is bounded, so by Theorem 4.1, passing to a subsequence we can assume that the C_i are all homotopic, and converge to a curve C .

By continuity, we now have $\ell_\phi(C) = |Z_\phi(\beta)|$. This implies that C is a union of saddle trajectories for ϕ , that is, a union of the γ_i . But as we just argued, the unique ϕ_i -geodesic representative in this homotopy class is the corresponding union of γ_i , and hence can only be a saddle connection if it is one of the γ_i . \square

5.9. Juggles. We conclude this section with a few brief remarks about the relationship between the ring-shrinking operation of the last few sections and the notion of a ‘juggle’ appearing in Gaiotto-Moore-Neitzke’s paper [14]. This material will not be used later and can be safely skipped.

Suppose that $\phi_+ \in \operatorname{Quad}(g, m)$ has a non-degenerate ring domain A . The closed trajectories of A have a certain hat-homology class $\alpha \in \hat{H}(\phi)$. Let $\delta \in \hat{H}(\phi)$ be the class of a saddle connection in A joining zeroes of ϕ_+ lying on different boundary components. Considering lines of suitable rational slope in the universal cover of A shows that for all $k \in \mathbb{Z}$ there are saddle connections for ϕ_+ with hat-homology class $\delta + k\alpha$. In particular, the spectrum $\Theta_{\phi_+} \subset S^1$ of phases θ for which $e^{i\pi\theta} \cdot \phi_+$ has a saddle trajectory has an accumulation point at $\theta = 0$.

By taking differentials $e^{i\pi\theta} \cdot \phi_+$ with θ varying near 0 we can define a path in $\operatorname{Quad}(g, m)$ with saddle-free endpoints which crosses infinitely many of the real codimension one walls that are the local connected components of F_2 . We refer to such a path as a juggle path. In Section 10.1 we will associate ideal triangulations to saddle-free differentials; the triangulations associated to the end-points of our path will then be related by a particular kind of infinite composition of flips, referred to in [14] as a juggle. The ring-shrinking move of Proposition 5.7 has the effect of removing the accumulation point at

$\theta = 0$ in the spectrum $\Theta_\phi \subset S^1$. This allows us to replace certain juggle paths by paths which meet only finitely many walls.





FIGURE 23. Local perturbations of a differential with two saddle trajectories of equal phase; the left perturbation has a ring domain, the right does not.

Let us consider the case when the boundary components of A are both closed saddle trajectories, and the differential ϕ_+ contains no other finite-length trajectories. After shrinking we obtain a differential ϕ with a closed curve made up of two saddle trajectories γ_1, γ_2 with hat-homology classes α_1, α_2 satisfying $\alpha = \alpha_1 + \alpha_2$. The trajectory structure of differentials near ϕ satisfying $\text{Im } Z_\phi(\alpha) = 0$ is determined by the wall

$$\text{Im } Z_\phi(\gamma_1)/Z_\phi(\gamma_2) = 0.$$

Differentials on the ϕ_+ side of this wall have a ring domain; differentials on the other side are saddle-free. The relevant geometry is illustrated in Figure 23.

The representation theory relevant to juggles is that of the Kronecker quiver (see also Example 12.5).



Let \mathcal{A} be the category of representations of this quiver, and let S_1, S_2 be the vertex simple objects, appropriately ordered. Stability conditions on \mathcal{A} satisfying

$$\text{Im } Z(S_1)/Z(S_2) > 0$$

have unique stable objects of dimension vectors $(n, n+1)$ and $(n+1, n)$ for all $n \geq 0$, and also a moduli space of stable objects of dimension vector $(1, 1)$ which is isomorphic to \mathbb{P}^1 . In particular, the set of phases of stable objects has an accumulation point. On the other hand, if

$$\text{Im } Z(S_1)/Z(S_2) < 0$$

then the only stable objects are the objects S_i themselves. The operation of ring-shrinking is the analogue of moving from a stability condition with $\text{Im } Z(S_1)/Z(S_2) > 0$ to one where $\text{Im } Z(S_1)/Z(S_2) = 0$. This has the effect of removing the accumulation point in the spectrum of stable phases.

6. COLLIDING ZEROES AND POLES: THE SPACES $\text{Quad}(\mathbb{S}, \mathbb{M})$

The spaces of quadratic differentials appearing in our main Theorems do not have fixed polar type; rather the zeroes are allowed to collide with the double poles. This means that we are dealing with spaces which are unions of strata of the form $\text{Quad}(g, m)$. It is convenient to label these spaces by diffeomorphism classes of marked bordered surfaces. For definitions concerning such surfaces see the Introduction or Section 8.1 below.

6.1. Union of strata. A GMN differential ϕ on a compact Riemann surface S determines a marked bordered surface (\mathbb{S}, \mathbb{M}) by the following construction. To define the surface \mathbb{S} we take the underlying smooth surface of S and perform an oriented real blow-up at each pole of ϕ of order > 2 . The marked points \mathbb{M} are then the poles of ϕ of order ≤ 2 , considered as points of the interior of \mathbb{S} , together with the points on the boundary of \mathbb{S} corresponding to the distinguished tangent directions of Section 3.3.

By a quadratic differential on a marked bordered surface (\mathbb{S}, \mathbb{M}) we mean a pair (S, ϕ) , consisting of a compact Riemann surface S and a GMN differential ϕ , whose associated marked bordered surface is diffeomorphic to (\mathbb{S}, \mathbb{M}) . We let $\text{Quad}(\mathbb{S}, \mathbb{M})$ denote the space of equivalence classes of such pairs.

A marked bordered surface (\mathbb{S}, \mathbb{M}) is determined up to diffeomorphism by the genus of \mathbb{S} , the number of punctures, and an unordered collection of positive integers encoding the number of marked points on each boundary component. In more concrete terms then, we have

$$(6.1) \quad \text{Quad}(\mathbb{S}, \mathbb{M}) = \bigcup_{(g, m)} \text{Quad}(g, m),$$

where the union is over pairs (g, m) , where $g = g(\mathbb{S})$ is the genus of \mathbb{S} , and there is one $m_i \in \{1, 2\}$ for each puncture $p \in \mathbb{P}$, and one $m_i = k_i + 2$ for each boundary component containing k_i marked points.

Let (g, m) be the unique pair appearing in the decomposition (6.1) for which all $m_i \geq 2$. In the proof of Proposition 2.4 we considered a vector bundle

$$(6.2) \quad \mathcal{H}(g, m) \longrightarrow \mathcal{M}(g, d) / \text{Sym}(m),$$

whose fibre over a marked curve $(S, (p_i))$ is the space of global sections of the line bundle $\omega_S^{\otimes 2}(\sum_i m_i p_i)$. The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ is the open subset of $\mathcal{H}(g, m)$ consisting of sections with simple zeroes which are disjoint from the points p_i for which $m_i > 2$. As in the proof of Proposition 2.4 it is therefore either empty, or a complex orbifold of dimension n .

Recall that a GMN differential is called *complete* if it has no simple poles.

Lemma 6.1. *The subset of complete differentials is an open subset*

$$(6.3) \quad \text{Quad}(\mathbb{S}, \mathbb{M})_0 = \text{Quad}(g, m) \subset \text{Quad}(\mathbb{S}, \mathbb{M})$$

whose complement is a normal crossings divisor.

Proof. Locally on the universal curve over $\mathcal{M}(g, d)$ we can trivialise the line bundle $\omega_S^{\otimes 2}(\sum_i m_i p_i)$. Working locally on $\mathcal{H}(g, m)$ we can therefore associate to each point an unordered collection of complex numbers $\{r_p : p \in \mathbb{P}\}$ obtained by evaluating the defining section ϕ at the marked points p_i for which $m_i \leq 2$. The resulting locally-defined functions r_p are holomorphic on $\mathcal{H}(g, m)$, and the complement of the open stratum (6.3) is precisely the vanishing locus of the product of these functions.

Suppose that a point $\phi \in \text{Quad}(\mathbb{S}, \mathbb{M})$ has $s \geq 1$ simple poles. Then the locally-defined map

$$r : \text{Quad}(\mathbb{S}, \mathbb{M}) \rightarrow \mathbb{C}^s$$

given by the functions r_p corresponding to the simple poles of ϕ is a submersion at ϕ . Indeed, using Riemann-Roch, for each simple pole p of ϕ we can find sections of $\omega_S^{\otimes 2}(\sum_i m_i p_i)$ which vanish at all the other simple poles of ϕ but not at p . Adding linear combinations of such sections to ϕ shows that r has a locally-defined section. It follows from this that the complement of the open stratum (6.3) is a normal crossings divisor. \square

6.2. Signed differentials. Fix a marked bordered surface (\mathbb{S}, \mathbb{M}) . Although the hat-homology groups $\hat{H}(\phi)$ form a local system over the orbifold $\text{Quad}(\mathbb{S}, \mathbb{M})_0$, this is not true over the larger orbifold $\text{Quad}(\mathbb{S}, \mathbb{M})$, since by Lemma 2.2, at differentials where a zero has collided with a double pole the rank of the hat-homology group drops by one.

A stronger statement is that the local system of hat-homology groups over $\text{Quad}(\mathbb{S}, \mathbb{M})_0$ cannot be extended to a local system on $\text{Quad}(\mathbb{S}, \mathbb{M})$. The reason is that parallel transport around a differential with a simple pole at a point p changes the sign of the residue class β_p (see the proof of Lemma 6.2 below).

By a *signed quadratic differential* on (\mathbb{S}, \mathbb{M}) we mean a differential

$$(S, \phi) \in \text{Quad}(\mathbb{S}, \mathbb{M})$$

together with a choice of sign of the residue $\text{Res}_p(\phi)$ at each puncture $p \in \mathbb{P}$. Note that by (3.1) this is equivalent to choosing a square-root of the function r_p of the last paragraph. The set of such signed differentials therefore forms a smooth complex orbifold equipped with a finite map

$$\text{Quad}^\pm(\mathbb{S}, \mathbb{M}) \rightarrow \text{Quad}(\mathbb{S}, \mathbb{M})$$

branched precisely over the complement of the incomplete locus. We write $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})_0$ for the open subset of $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})$ consisting of signed differentials whose underlying differential is complete.




FIGURE 24. The local monodromy as a zero encircles a double pole.

Lemma 6.2. *The local system of hat-homology groups $\hat{H}(\phi)$ pulled back to the étale cover $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})_0 \rightarrow \text{Quad}(\mathbb{S}, \mathbb{M})_0$ extends to a local system on $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})$.*

Proof. We must compute the monodromy of the hat-homology local system around each component of the boundary divisor consisting of non-complete differentials. Consider a differential ϕ_0 lying on this divisor, having a single simple pole p_0 . Nearby complete differentials ϕ will have a corresponding double pole p and a simple zero q which have collided to produce p_0 .

The hat-homology group of ϕ is spanned by the hat-homology classes of saddle connections. Saddle connections of ϕ not ending at q correspond canonically to saddle connections of ϕ_0 not ending at p_0 , and their hat-homology classes are therefore unaffected by the local monodromy around ϕ_0 . Consider the class α_q of a saddle connection ending at q , and let β_p be the residue class at p . The local monodromy of the Gauss-Manin connection⁶ acts on the classes (α_q, β_p) by the transformation

$$(6.4) \quad \alpha_q \mapsto \alpha_q + \beta_p, \quad \beta_p \mapsto -\beta_p,$$

see Figure 24. This transformation has order 2 and hence becomes trivial when pulled-back to the double cover determined by a choice of sign of $Z_\phi(\beta_p)$. \square

6.3. Extended hat-homology group. Let us consider the quotient orbifold

$$(6.5) \quad \text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) = \text{Quad}^\pm(\mathbb{S}, \mathbb{M}) / \mathbb{Z}_2^{\oplus \mathbb{P}},$$

where $\mathbb{Z}_2^{\oplus \mathbb{P}}$ acts in the obvious way on the signings. Note that this quotient is to be understood in the category of spaces over $\text{Quad}(\mathbb{S}, \mathbb{M})$, since the punctures \mathbb{P} form a non-trivial local system over this space. Practically speaking, we can locally trivialise this local system \mathbb{P} on $\text{Quad}(\mathbb{S}, \mathbb{M})$, define local quotients by the group $\mathbb{Z}_2^{\oplus \mathbb{P}}$, and then glue these together to form the global quotient (6.5). Note that there is an open inclusion

$$\text{Quad}(\mathbb{S}, \mathbb{M})_0 \subset \text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}).$$

The only difference between the spaces $\text{Quad}(\mathbb{S}, \mathbb{M})$ and $\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$ is some extra orbifolding along the incomplete locus.

⁶To check the sign change for the residue class β_p , it may be helpful to consider the family of differentials $(z - a)dz^{\otimes 2}/z^2$ of residue $4\pi i\sqrt{a}$, as a encircles the origin.

The local system of Lemma 6.2 descends to the orbifold $\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$. The *extended hat-homology group* $\hat{H}^e(\phi)$ of a GMN differential ϕ is defined to be the fibre of this local system at ϕ . This group coincides with the usual hat-homology group $\hat{H}(\phi)$ precisely if ϕ is complete. In general $\hat{H}^e(\phi)$ comes equipped with a skew-symmetric pairing and canonically defined residue classes β_p , one for each simple or even order pole of ϕ . This data is obtained by parallel transport from a nearby complete differential.

Lemma 6.3. *For any GMN differential ϕ_0 there is a canonical group homomorphism*

$$q: \hat{H}^e(\phi_0) \rightarrow \hat{H}(\phi_0)$$

whose kernel is spanned over \mathbb{Q} by the residue classes β_p corresponding to the simple poles of ϕ_0 .

Proof. Consider the family of spectral covers $\hat{S} \rightarrow S$ defined by differentials ϕ in some small neighbourhood $\phi_0 \in U \subset \text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$. These covers vary holomorphically because the divisor E of formula (2.3) varies holomorphically, and φ is a holomorphically varying section. Note however that the open surface \hat{S}° changes discontinuously in general, as it must, since the rank of the hat-homology group drops at differentials with simple poles.

More precisely, when a zero of the differential ϕ collides with a double pole p , the infinite critical point p becomes a finite critical point of ϕ_0 which is moreover a branch-point of the corresponding spectral cover. Thus the two punctures in the surface \hat{S} lying over p which are removed when defining \hat{S}° collide and get filled in.

Define a subsurface $\hat{S}' \subset \hat{S}$ by removing from \hat{S} the inverse images of those infinite critical points of ϕ which remain infinite critical points for the differential ϕ_0 . The homology groups $H_1(\hat{S}'; \mathbb{Z})^-$ form a local system over U whose fibre at ϕ_0 coincides with $H_1(\hat{S}^\circ; \mathbb{Z})^-$. Over the complete locus, the inclusion $\hat{S}^\circ \subset \hat{S}'$ defines a map of local systems

$$q: H_1(\hat{S}^\circ; \mathbb{Z})^- \rightarrow H_1(\hat{S}'; \mathbb{Z})^-.$$

The same analysis we used to prove Lemma 2.3 shows that the kernel of q is spanned over \mathbb{Q} by the residue classes β_p corresponding to the simple poles of ϕ . Specialising the map q to the fibres at ϕ_0 then gives the result. \square

6.4. Blowing up simple poles. In this section we explain a surgery which, whilst not required in the proofs of the main theorems of the paper, helps explain the geometry as zeroes collide with double poles and one passes between different strata in $\text{Quad}(\mathbb{S}, \mathbb{M})$.

The surgery involves ‘blowing up’ a simple pole and inserting a metric cylinder (i.e. a disk with differential $r dz^{\otimes 2}/z^2$ for some $r \in \mathbb{R}_{<0}$). Although as topological surfaces the complement of the inserted cylinder differs from the original surface by a real blow-up,

metrically the surfaces are related by slitting a finite length of trajectory and opening up the slit into a boundary component, as indicated in Figure 25.




FIGURE 25. Replacing a simple pole with a degenerate ring domain.

Proposition 6.4. *Let ϕ be a GMN differential on S with s simple poles p_i . Let $r_i \in \mathbb{R}_{>0}$ be sufficiently small. Then there is a uniquely-defined complete GMN differential with double poles at the p_i , centred on degenerate ring domains with parameters $-r_i$, and equivalent on the complement of the closures of those ring domains to $(S \setminus \cup_i \gamma_i, \phi)$, where γ_i is the unique horizontal trajectory of ϕ of length r_i with one end-point at p_i .*

Proof. If the r_i are sufficiently small, the trajectory γ_i is embedded in the surface and does not contain any critical points other than p_i . The existence part of the statement is then depicted in Figure 25. The dashed rectangle has boundaries on the horizontal and vertical foliations for the relevant differentials; it lies in a co-ordinate chart in the central picture, and is mapped conformally in its interior to the two outer pictures, which define surfaces with quadratic differentials which are equivalent in the component exterior to the rectangle's boundary arc. The existence of the differential on the surface on the right, obtained by gluing a cylinder and a surface with geodesic boundary containing a simple zero, is an application of Lemma 3.3. \square

If the original differential has finite area, there are closed geodesics for a dense set of phases [33, Theorem 25.2]; however, the only such geodesics which survive to the surgered surface are those which are disjoint from the length r_i segments of the trajectories emanating from the simple poles. This is compatible with the fact that the spectrum of phases of closed geodesics after surgery is closed in S^1 .

Remark 6.5. *Let us use the notation $\text{Quad}(g, (1^d))$ for the space of GMN differentials (S, ϕ) with S of genus g , and ϕ having d simple poles. Similarly let $\text{Quad}_r(g, (2^d))$ denote the space of GMN differentials (S, ϕ) with S of genus g and ϕ having d double poles, each of residue $\pm r$. The construction of Proposition 6.4 gives an injective map*

$$B: \text{Quad}_r(g, (2^d)) \rightarrow \text{Quad}(g, (1^d))$$

which moreover commutes with the locally-defined period maps on both sides. It is not hard to convince oneself that B is in fact a local homeomorphism. This reduces the

question of whether the period map is a local isomorphism to the case of differentials with at least one infinite critical point.

A closely related model is obtained by opening up a length l segment of an irrational foliation on a torus $S = T^2$ to obtain a recurrent surface with one boundary component, the boundary made of two equal length saddle trajectories; one can then glue in a degenerate ring domain centred on a double pole as above to obtain another foliation on a closed torus. The closed geodesics in the new (infinite area) surface correspond to the (p, q) -curves on T^2 which are disjoint from the original straight arc of length l . It is easy to check that only finitely many (p, q) -curves have this property, hence the surgery collapses the spectrum of closed trajectories from a dense subset of the circle to a finite subset.

6.5. Extended period map. Let (S, \mathbb{M}) be a marked bordered surface, and let

$$\text{Quad}(g, m) = \text{Quad}(S, \mathbb{M})_0 \subset \text{Quad}(S, \mathbb{M})$$

be the corresponding open stratum of complete differentials. Fix a free abelian group Γ of rank n given by (2.2). By an *extended framing* of a point $\phi \in \text{Quad}_\heartsuit(S, \mathbb{M})$ we mean an isomorphism of groups

$$\theta: \Gamma \rightarrow \hat{H}^e(\phi).$$

Defining the space $\text{Quad}^\Gamma(S, \mathbb{M})$ of extended framed differentials in the obvious way, we obtain an unbranched cover

$$\text{Quad}^\Gamma(S, \mathbb{M}) \rightarrow \text{Quad}_\heartsuit(S, \mathbb{M}).$$

Over the locus $\text{Quad}(S, \mathbb{M})_0$ of complete differentials, the resulting space coincides with the space $\text{Quad}^\Gamma(g, m)$ considered before.

Lemma 2.6 shows that the generic automorphism group of the orbifold $\text{Quad}(S, \mathbb{M})$ is trivial except when (S, \mathbb{M}) is one of

- (i) an unpunctured disc with 3 or 4 points on its boundary;
- (ii) an annulus with one marked point on each boundary component;
- (iii) a closed torus with a single puncture;

corresponding to polar types (5), (6) and (3, 3) in genus $g = 0$, and polar type (2) in genus $g = 1$. As explained before, in all these cases the orbifold $\text{Quad}(S, \mathbb{M})$ also has a non-trivial generic automorphism group.

Proposition 6.6. *Assume that (S, \mathbb{M}) is not one of the 4 exceptional surfaces listed above. Then the space $\text{Quad}^\Gamma(S, \mathbb{M})$ is a complex manifold. The period map extends to a local isomorphism of complex manifolds*

$$\pi: \text{Quad}^\Gamma(S, \mathbb{M}) \rightarrow \text{Hom}_{\mathbb{Z}}(\Gamma, \mathbb{C}).$$

Proof. Assume first that (\mathbb{S}, \mathbb{M}) is not a sphere with 3 or 4 punctures. Suppose that a point of $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$ has a non-trivial automorphism. This means that the underlying differential ϕ has a non-trivial automorphism which acts trivially on the extended hat-homology group $\hat{H}^e(\phi)$. It follows from Lemma 6.3 that this automorphism also acts trivially on the hat-homology group $\hat{H}(\phi)$. But we proved in Theorem 4.12 that no such automorphisms exist. Thus $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$ is a manifold.

The extended period map π is defined in the obvious way: the period of a differential ϕ defines a map $Z_\phi: \hat{H}(\phi) \rightarrow \mathbb{C}$ which induces a group homomorphism $Z_\phi: \hat{H}^e(\phi) \rightarrow \mathbb{C}$ by composing with the map q of Lemma 6.3. To show that π is a local isomorphism, suppose that a nonzero tangent vector v to $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$ at some point ϕ lies in the kernel of the derivative of π . Then, since the strata of $\text{Quad}(\mathbb{S}, \mathbb{M})$ are determined by the vanishing of the periods $Z_\phi(\beta_p)$, it follows that v is tangent to the stratum containing ϕ . But the period map is a local isomorphism on each stratum by Theorem 4.12, so this gives a contradiction.

In the case when (\mathbb{S}, \mathbb{M}) is a 3 or 4 punctured sphere, the above proof is incomplete, because the two polar types $(1, 1, 2)$ and $(1, 1, 1, 1)$ in genus $g = 0$ were excluded from Theorem 4.12, since the corresponding spaces $\text{Quad}(g, m)$ have a non-trivial generic automorphism group. In both of these cases the generic automorphisms identified in Example 2.7 act non-trivially on the extended hat-homology group, since they permute the simple poles, and hence the corresponding residue classes β_p . The fact that the period map is a local isomorphism on the corresponding strata of $\text{Quad}(\mathbb{S}, \mathbb{M})$ can be proved exactly as in Theorem 4.12, or just checked directly. \square

6.6. Degenerations. We finish this first part of the paper with two technical results which will be used later in the proofs of our main Theorems. The first one will allow us to extend our correspondence between differentials and stability conditions over the incomplete locus in $\text{Quad}(\mathbb{S}, \mathbb{M})$.

Proposition 6.7. *Take a framed differential $\phi_0 \in \text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$. Then for any $\epsilon > 0$ there is a neighbourhood $\phi_0 \in U \subset \text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$ such that for any differential $\phi \in U$, and any class $\gamma \in \Gamma$ represented by a non-closed saddle connection in ϕ , there is an inequality*

$$|Z_\phi(\gamma) - Z_{\phi_0}(\gamma)| < \epsilon |Z_{\phi_0}(\gamma)|.$$

Proof. We can assume that all differentials $\phi \in U$ are on a fixed underlying smooth surface \mathbb{S} , with finite critical points at fixed points $x_i \in \mathbb{S}$. However we must allow the double poles of ϕ to move, so that they can collide with the zeroes. We can assume that if ϕ has a simple pole at x_i then so does ϕ_0 .

Consider the subset of $U \times \mathbb{S}$ consisting of pairs (ϕ, y) with $y \in \mathbb{S}$ lying on a non-closed saddle connection for ϕ , and let F be its closure. Then F contains no points of the form

(ϕ, p_j) with p_j an infinite critical point of ϕ , because any such point is contained in a trapping neighbourhood containing no non-closed saddle connections. Thus, shrinking U if necessary, we have a bound

$$|\sqrt{\phi} - \sqrt{\phi_0}| < \epsilon |\sqrt{\phi_0}|$$

for all points of F . Integrating this along a non-closed saddle connection for ϕ gives the result. \square

The next result is a kind of completeness result for the space $\text{Quad}(\mathbb{S}, \mathbb{M})$. It will be used later to prove that the image of the map we construct from differentials to stability conditions is closed. We say that a saddle connection γ for a GMN differential ϕ is *degenerate* if it is closed, and is moreover freely homotopic in S° to a small loop around a double pole of ϕ .⁷

Proposition 6.8. *Consider a sequence of framed, complete differentials*

$$\phi_n \in \text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})_0, \quad n \geq 1,$$

whose periods $Z_{\phi_n} : \Gamma \rightarrow \mathbb{C}$ converge. Suppose moreover that there is a universal constant $L > 0$ such that any non-degenerate saddle connection for ϕ_n has length $\geq L$. Then some subsequence of the points ϕ_n converges to a limit in $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$.

Proof. Using the fact that stable curves are Gorenstein it is easy to see that the vector bundle (6.2) extends to a bundle

$$\bar{\mathcal{H}}(g, m) \longrightarrow \bar{\mathcal{M}}(g, d) / \text{Sym}(m)$$

over the Deligne-Mumford compactification. The projectivisation of this bundle is a compact space, and so, passing to a subsequence, and potentially rescaling the ϕ_n , we can assume that the differentials ϕ_n have a limit $(S, \phi) \in \bar{\mathcal{H}}(g, m)$. The hypothesis that the periods Z_{ϕ_n} converge then implies that the rescaling must have been unnecessary.

Note that our assumption implies that if γ is a path in the surface S_n which either connects two finite critical points of ϕ_n , or is closed and not homotopic to a small loop around a double pole, then the length of γ in the ϕ_n -metric is at least L . Indeed, by Theorem 4.1, the curve γ has a minimal geodesic in its homotopy class and this is a union of saddle connections.

Suppose that the limit curve S has a node p , and that the limit section ϕ is non-vanishing at p . Note that the induced quadratic differential on the normalization has a double pole at the inverse image of p . Consider a curve connecting two zeroes on S_n lying on opposite sides of the neck which shrinks to the node p . Then as $n \rightarrow \infty$ the period of the corresponding hat-homology class diverges to infinity, which contradicts the fact that the periods Z_{ϕ_n} converge.

⁷ Note that a saddle connection γ of phase θ is nothing but a saddle trajectory for the rotated differential $e^{-i\pi\theta} \cdot \phi$, and that γ is degenerate precisely if this saddle trajectory is the boundary of a degenerate ring domain.

Suppose instead that S is a stable curve with a node p , and that the section ϕ vanishes at p . Then consider a closed curve γ on S_n encircling the neck, homotopic to the vanishing cycle. Either γ is non-separating, or there is more than one marked point on each side of γ , so γ cannot be homotopic to a small loop around a double pole. Then as $n \rightarrow \infty$ the length in the ϕ_n -metric of γ tends to zero, which again gives a contradiction.

Thus we conclude that the limit curve S is non-singular. Suppose that the limit differential ϕ is defined by a section of $\omega_S^{\otimes 2}(\sum_i m_i p_i)$ which has a zero of some order $k \geq 1$ at a point $p \in S$. This means that k simple zeroes of the ϕ_n have collided in the limit. If $p = p_i$ is a marked point then we set $m = m_i$, and otherwise we set $m = 0$.

Suppose first that $k - m \leq -2$ so that p is an infinite critical point of ϕ . Take a path γ_n which connects two zeroes z_1, z_2 of ϕ_n , both of which tend to p . Then the period of the corresponding hat-homology class tends to infinity, contradicting the assumptions. If $k = 1$ so that there is only one zero z which tends to p , then we must have $m \geq 3$, and we can take γ_n to be a small loop around p based at z .

On the other hand, if $k - m > -2$ then p is a finite critical point of ϕ . If also $k \geq 2$, then take a path γ_n connecting two zeroes of ϕ_n which both tend to p . The length of this path will tend to zero as $n \rightarrow \infty$ which gives a contradiction as before. We conclude that ϕ has simple zeroes distinct from the points p_i of order $m_i > 2$. This is precisely the condition that $\phi \in \text{Quad}(\mathbb{S}, \mathbb{M})$. \square

7. QUIVERS AND STABILITY CONDITIONS

This section consists of fairly well-known background material on t-structures and tilting, quivers with potential and their mutations, and stability conditions.

7.1. Introduction. Let \mathcal{D} be a k -linear triangulated category of finite type. We denote the shift functor by $[1]$ and use the notation

$$\text{Hom}_{\mathcal{D}}^i(A, B) := \text{Hom}_{\mathcal{D}}(A, B[i]).$$

The finite type condition is the statement that for all objects $A, B \in \mathcal{D}$

$$\dim_k \bigoplus_{i \in \mathbb{Z}} \text{Hom}_{\mathcal{D}}^i(A, B) < \infty.$$

The Grothendieck group $K(\mathcal{D})$ then carries the Euler bilinear form

$$\chi(-, -): K(\mathcal{D}) \times K(\mathcal{D}) \rightarrow \mathbb{Z}$$

defined by the formula

$$\chi(E, F) = \sum_{i \in \mathbb{Z}} (-1)^i \dim_k \text{Hom}_{\mathcal{D}}^i(E, F).$$

Beginning in Section 9 we shall focus on the particular properties of the categories $\mathcal{D} = \mathcal{D}(\mathbb{S}, \mathbb{M})$ appearing in our main Theorems, but the present section consists of general theory, and the only properties of \mathcal{D} that will be important are

- (i) \mathcal{D} admits a bounded t-structure whose heart $\mathcal{A} \subset \mathcal{D}$ is of finite length and has a finite number n of simple objects up to isomorphism;
- (ii) \mathcal{D} is a CY_3 category, meaning that there are functorial isomorphisms

$$\text{Hom}_{\mathcal{D}}^i(A, B) \cong \text{Hom}_{\mathcal{D}}^{3-i}(B, A)^* \text{ for all objects } A, B \in \mathcal{D}.$$

Note that (i) implies that $K(\mathcal{D}) \cong \mathbb{Z}^{\oplus n}$ is free of finite rank, and (ii) implies that the Euler form is skew-symmetric.

The main point of this section is to expand on the following two statements:

- (a) Associated to any triangulated category \mathcal{D} there is a complex manifold $\text{Stab}(\mathcal{D})$ of dimension n parameterizing certain structures on \mathcal{D} known as stability conditions. When \mathcal{D} satisfies condition (i), a large open subset of $\text{Stab}(\mathcal{D})$ can be described as a union of cells, one for each bounded t-structure with finite-length heart. The way these cells are glued together along their boundaries is controlled by an abstract operation called tilting.
- (b) A large class of triangulated categories \mathcal{D} satisfying both conditions (i) and (ii) can be defined using quivers with potential via the Ginzburg algebra construction. The abstract tilting operation referred to in (a) can then be described concretely in terms of mutations of quivers with potential.

7.2. Hearts and tilting. Let \mathcal{D} be a triangulated category. We shall be concerned with bounded t-structures on \mathcal{D} . Any such t-structure is determined by its heart $\mathcal{A} \subset \mathcal{D}$, which is a full abelian subcategory. We use the term *heart* to mean the heart of a bounded t-structure. A heart will be called *finite-length* if it is artinian and noetherian as an abelian category.

We say that a pair of hearts $(\mathcal{A}_1, \mathcal{A}_2)$ in \mathcal{D} is a *tilting pair* if the equivalent conditions

$$\mathcal{A}_2 \subset \langle \mathcal{A}_1, \mathcal{A}_1[-1] \rangle, \quad \mathcal{A}_1 \subset \langle \mathcal{A}_2[1], \mathcal{A}_2 \rangle$$

are satisfied.⁸ We also say that \mathcal{A}_1 is a *left tilt* of \mathcal{A}_2 , and that \mathcal{A}_2 is a *right tilt* of \mathcal{A}_1 . Note that $(\mathcal{A}_1, \mathcal{A}_2)$ is a tilting pair precisely if so is $(\mathcal{A}_2[1], \mathcal{A}_1)$.

If $(\mathcal{A}_1, \mathcal{A}_2)$ is a tilting pair in \mathcal{D} , then the subcategories

$$\mathcal{T} = \mathcal{A}_1 \cap \mathcal{A}_2[1], \quad \mathcal{F} = \mathcal{A}_1 \cap \mathcal{A}_2$$

⁸The angular brackets here signify the *extension-closure* operation: given full subcategories $\mathcal{A}, \mathcal{B} \subset \mathcal{D}$, the extension-closure $\mathcal{C} = \langle \mathcal{A}, \mathcal{B} \rangle \subset \mathcal{D}$ is the smallest full subcategory of \mathcal{D} containing both \mathcal{A} and \mathcal{B} , and such that if $X \rightarrow Y \rightarrow Z \rightarrow X[1]$ is a triangle in \mathcal{D} with $X, Z \in \mathcal{C}$ then $Y \in \mathcal{C}$ also.




FIGURE 26. A tilting pair.




FIGURE 27. Left and right tilts of a heart.

form a torsion pair $(\mathcal{T}, \mathcal{F}) \subset \mathcal{A}_1$. Conversely, if $(\mathcal{T}, \mathcal{F}) \subset \mathcal{A}_1$ is a torsion pair, then the subcategory $\mathcal{A}_2 = \langle \mathcal{F}, \mathcal{T}[-1] \rangle$ is a heart, and the pair $(\mathcal{A}_1, \mathcal{A}_2)$ is a tilting pair.

A special case of the tilting construction will be particularly important. Suppose that \mathcal{A} is a finite-length heart and $S \in \mathcal{A}$ is a simple object. Let $\langle S \rangle \subset \mathcal{A}$ be the full subcategory consisting of objects $E \in \mathcal{A}$ all of whose simple factors are isomorphic to S . Define full subcategories

$$S^\perp = \{E \in \mathcal{A} : \text{Hom}_{\mathcal{A}}(S, E) = 0\}, \quad {}^\perp S = \{E \in \mathcal{A} : \text{Hom}_{\mathcal{A}}(E, S) = 0\}.$$

One can either view $\langle S \rangle$ as the torsion part of a torsion pair on \mathcal{A} , in which case the torsion-free part is S^\perp , or as the torsion-free part, in which case the torsion part is ${}^\perp S$. We can then define tilted hearts

$$\mu_S^-(\mathcal{A}) = \langle S[1], {}^\perp S \rangle, \quad \mu_S^+(\mathcal{A}) = \langle S^\perp, S[-1] \rangle,$$

which we refer to as the left and right tilts of the heart \mathcal{A} at the simple S . They fit into tilting pairs $(\mu_S^-(\mathcal{A}), \mathcal{A})$ and $(\mathcal{A}, \mu_S^+(\mathcal{A}))$. Note the relation

$$(7.1) \quad \mu_{S[1]}^+ \circ \mu_S^-(\mathcal{A}) = \mathcal{A}.$$

The *tilting graph* of \mathcal{D} is the graph $\text{Tilt}(\mathcal{D})$ whose vertices are finite-length hearts, and in which two vertices are joined by an edge if the corresponding hearts are related by a tilt in a simple object. There is a natural action of the group of triangulated autoequivalences $\text{Aut}(\mathcal{D})$ on this graph.

If $\mathcal{A} \subset \mathcal{D}$ is a finite-length heart we denote by $\text{Tilt}_{\mathcal{A}}(\mathcal{D}) \subset \text{Tilt}(\mathcal{D})$ the connected component containing \mathcal{A} . We say that the hearts in $\text{Tilt}_{\mathcal{A}}(\mathcal{D})$ are *reachable* from \mathcal{A} . We say that an autoequivalence $\Phi \in \text{Aut}(\mathcal{D})$ is *reachable* from \mathcal{A} if its action on $\text{Tilt}(\mathcal{D})$ preserves the connected component $\text{Tilt}_{\mathcal{A}}(\mathcal{D})$. These autoequivalences form a subgroup $\text{Aut}_{\mathcal{A}}(\mathcal{D}) \subset \text{Aut}(\mathcal{D})$.

We say that a finite-length heart $\mathcal{A} \subset \mathcal{D}$ is *infinitely tiltable* if the graph $\text{Tilt}_{\mathcal{A}}(\mathcal{D})$ is $2n$ -regular, where n is the rank of $K(\mathcal{D})$. This means that the tilting process can be continued indefinitely at all simple objects, and in both directions, without leaving the class of finite-length hearts.

7.3. Tilting in the CY_3 case. Suppose now that \mathcal{D} is a triangulated category with the CY_3 property. To ensure the existence of the twist functors appearing below we should also assume that \mathcal{D} is algebraic in the sense of Keller [18, Section 3.6].

Associated to a finite-length heart $\mathcal{A} \subset \mathcal{D}$ there is a quiver $Q(\mathcal{A})$, whose vertices are indexed by the isomorphism classes of simple objects $S_i \in \mathcal{A}$ and which has

$$n_{ij} = \dim_k \text{Ext}_{\mathcal{A}}^1(S_i, S_j)$$

arrows connecting vertex i to vertex j . We call a finite-length heart $\mathcal{A} \subset \mathcal{D}$ *non-degenerate* if it is infinitely-tiltable and if, for every heart $\mathcal{B} \subset \mathcal{D}$ reachable from \mathcal{A} , the quiver $Q(\mathcal{B})$ has no loops or oriented 2-cycles.

The quiver $Q(\mathcal{A})$ associated to a finite-length heart $\mathcal{A} \subset \mathcal{D}$ has no loops precisely if the simple objects of \mathcal{A} are all spherical in the sense of [30]. Any spherical object $S \in \mathcal{D}$ defines an autoequivalence $\text{Tw}_S \in \text{Aut}(\mathcal{D})$ called a spherical twist. It has the property that for any object $E \in \mathcal{D}$ there is a triangle

$$\text{Hom}_{\mathcal{D}}^{\bullet}(S, E) \otimes S \longrightarrow E \longrightarrow \text{Tw}_S(E).$$

Suppose now that $\mathcal{A} \subset \mathcal{D}$ is a non-degenerate finite-length heart with n simple objects up to isomorphism. Taken together, the spherical twists in these simple objects generate a subgroup

$$\text{Sph}_{\mathcal{A}}(\mathcal{D}) = \langle \text{Tw}_{S_1}, \dots, \text{Tw}_{S_n} \rangle \subset \text{Aut}(\mathcal{D}).$$

The following result is well-known, but for the reader's convenience we include a sketch proof. A more careful treatment can be found for example in [22].

Proposition 7.1. (a) *for every simple object $S \in \mathcal{A}$ there is a relation*

$$\text{Tw}_S(\mu_S^-(\mathcal{A})) = \mu_S^+(\mathcal{A}) \subset \mathcal{D};$$

(b) if \mathcal{B} is reachable from \mathcal{A} , then $\mathrm{Sph}_{\mathcal{A}}(\mathcal{D}) = \mathrm{Sph}_{\mathcal{B}}(\mathcal{D})$.

Proof. Take a simple object $S \in \mathcal{A}$. Since $Q(\mathcal{A})$ has no oriented 2-cycles, we can order the simple objects of \mathcal{A} so that $S = S_i$, and $\mathrm{Ext}^1(S_j, S_i) = 0$ for $j < i$ and $\mathrm{Ext}^1(S_i, S_j) = 0$ for $j > i$. Then for all $j < i$, the object $\mathrm{Tw}_{S_i}(S_j)$ is the universal extension

$$0 \longrightarrow S_j \longrightarrow \mathrm{Tw}_{S_i}(S_j) \longrightarrow \mathrm{Ext}_{\mathcal{A}}^1(S_i, S_j) \otimes S_i \longrightarrow 0.$$

This clearly lies in $S_i^\perp \subset \mu_{S_i}^+(\mathcal{A})$, and is easily checked to be simple. In this way one sees that the simple objects of $\mu_{S_i}^+(\mathcal{A})$ are

$$(\mathrm{Tw}_{S_i}(S_1), \dots, \mathrm{Tw}_{S_i}(S_{i-1}), S_i[-1], S_{i+1}, \dots, S_n).$$

By a similar argument, or using $\mu_{S_i[-1]}^- \mu_{S_i}^+(\mathcal{A}) = \mathcal{A}$, it follows that the simple objects in $\mu_{S_i}^-(\mathcal{A})$ are

$$(S_1, \dots, S_{i-1}, S_i[1], \mathrm{Tw}_{S_i}^{-1}(S_{i+1}), \dots, \mathrm{Tw}_{S_i}^{-1}(S_n)).$$

Property (a) is then clear, using the identity $\mathrm{Tw}_{S_i}(S_i) = S_i[-2]$, and the fact that a finite-length heart is determined by its simple objects. Property (b) follows from the identity

$$\mathrm{Tw}_{\mathrm{Tw}_{S_i}(S_j)} = \mathrm{Tw}_{S_i} \circ \mathrm{Tw}_{S_j} \circ \mathrm{Tw}_{S_i}^{-1},$$

and the fact that $\mathrm{Tw}_{S_i[1]} = \mathrm{Tw}_{S_i}$. □

7.4. Quivers with potential. Suppose that \mathcal{D} is a triangulated category with the CY_3 property. In Section 7.3 we associated a quiver $Q(\mathcal{A})$ to a finite-length heart $\mathcal{A} \subset \mathcal{D}$ encoding the dimensions of the extension spaces between the simple objects. The next result shows that after making the choice of a potential on $Q(\mathcal{A})$ one can reverse this process.

For all notions regarding quivers with potential we refer to [8, §2–5] and [20, §2]. In particular, we recall that a *potential* on a quiver Q is a formal linear combination of oriented cycles in Q , and that a potential is called *reduced* if it is a sum of cycles of length ≥ 3 .

Theorem 7.2. *Associated to a quiver with reduced potential (Q, W) there is a CY_3 triangulated category $\mathcal{D}(Q, W)$ of finite type over k , with a bounded t -structure whose heart*

$$\mathcal{A} = \mathcal{A}(Q, W) \subset \mathcal{D}(Q, W)$$

is of finite-length and has associated quiver $Q(\mathcal{A})$ isomorphic to Q .

Proof. Define the category $\mathcal{D}(Q, W)$ to be the subcategory of the derived category of the complete Ginzburg algebra $\Pi(Q, W)$ consisting of objects with finite-dimensional cohomology. It has the CY_3 property by [20, Lemma 7.16, Theorem 7.17].

The category $\mathcal{D}(Q, W)$ has a standard bounded t-structure [20, Lemma 5.2] whose heart $\mathcal{A}(Q, W)$ is equivalent to the category of finite-dimensional modules for the complete Jacobi algebra $J(Q, W) = H^0(\Pi(Q, W))$. The algebra $J(Q, W)$ is the quotient of the complete path algebra of Q by the relations obtained by cyclically differentiating the potential W . In particular, its simple modules are naturally in bijection with the vertices of Q , and the spaces of extensions between them are based by the arrows in Q . \square

The combinatorial incarnation of the process of tilting at a simple module is called *mutation*. It acts on right-equivalence classes of quivers with potential. Roughly speaking, two potentials on a quiver Q are said to be *right-equivalent* if they differ by an automorphism of the completed path algebra which fixes the zero length paths; for the full definition see [8, §4] or [20, §2.1]. Right-equivalent potentials give rise to isomorphic complete Ginzburg algebras [20, Lemma 2.9], and hence equivalent categories $\mathcal{D}(Q, W)$.

Suppose that (Q, W) is a reduced quiver with potential, and fix a vertex i of Q . The mutation $(Q', W') = \mu_i(Q, W)$ is another reduced quiver with potential, well-defined up to right-equivalence, and depending only on the right-equivalence class of (Q, W) . The vertex sets of Q and Q' are naturally identified, and the operation μ_i is an involution. We refer the reader to [8, §5] or [20, §2.4] for the relevant definitions.

For our purposes, the importance of mutations of quivers with potential is the following result of Keller and Yang [20, Thm. 3.2, Cor. 5.5].

Theorem 7.3. *Let (Q, W) be a quiver with reduced potential, such that Q has no loops or oriented 2-cycles. Let i be a vertex of Q and set*

$$(Q', W') = \mu_i(Q, W).$$

Then there is a canonical pair of k -linear triangulated equivalences

$$\Phi_{\pm} : \mathcal{D}(Q', W') \longrightarrow \mathcal{D}(Q, W)$$

which induce tilts in the simple object $S_i \in \mathcal{A}(Q, W)$ in the sense that

$$\Phi_{\pm}(\mathcal{A}(Q', W')) = \mu_{S_i}^{\pm}(\mathcal{A}(Q, W)) \subset \mathcal{D}(Q, W),$$

and which moreover induce the natural bijections on simple objects.

For the last part of this statement, recall that there is a natural bijection between the simple objects of \mathcal{A} and those of $\mu_{S_i}^{\pm}(\mathcal{A})$ which was made explicit in the proof of Proposition 7.1. There is also a natural bijection between the vertices of the quivers (Q, W) and $\mu_i(Q, W)$. The claim is that these bijections are compatible with the canonical bijections between the vertices of the quivers Q, Q' and the simple objects in the corresponding standard hearts.

A quiver with potential (Q, W) is called *non-degenerate* [8, §7] if any sequence of mutations of (Q, W) results in a quiver with potential having no loops or oriented 2-cycles. Theorem 7.3 shows that this condition is equivalent to the statement that the standard heart $\mathcal{A}(Q, W) \subset \mathcal{D}(Q, W)$ is non-degenerate in the sense of Section 7.3.

7.5. Stability conditions. Here we summarize the basic properties of spaces of stability conditions. We refer the reader to [2, 3] for more details on this material. Let us fix a triangulated category \mathcal{D} , and assume for simplicity that the Grothendieck group $K(\mathcal{D}) \cong \mathbb{Z}^{\oplus n}$ is free of finite rank.

A *stability condition* $\sigma = (Z, \mathcal{P})$ on \mathcal{D} consists of a group homomorphism $Z: K(\mathcal{D}) \rightarrow \mathbb{C}$ called the *central charge*, and full additive subcategories $\mathcal{P}(\phi) \subset \mathcal{D}$ for each $\phi \in \mathbb{R}$, which together satisfy the following axioms:

- (a) if $E \in \mathcal{P}(\phi)$ then $Z(E) \in \mathbb{R}_{>0} \cdot e^{i\pi\phi} \subset \mathbb{C}$,
- (b) for all $\phi \in \mathbb{R}$, $\mathcal{P}(\phi + 1) = \mathcal{P}(\phi)[1]$,
- (c) if $\phi_1 > \phi_2$ and $A_j \in \mathcal{P}(\phi_j)$ then $\text{Hom}_{\mathcal{D}}(A_1, A_2) = 0$,
- (d) for each nonzero object $E \in \mathcal{D}$ there is a finite sequence of real numbers

$$\phi_1 > \phi_2 > \cdots > \phi_k$$

and a collection of triangles

$$0 = E_0 \xrightarrow{\quad} E_1 \xrightarrow{\quad} E_2 \longrightarrow \cdots \longrightarrow E_{k-1} \xrightarrow{\quad} E_k = E$$

with $A_j \in \mathcal{P}(\phi_j)$ for all j .

The semistable objects A_j appearing in the filtration of axiom (d) are unique up to isomorphism, and are called the *semistable factors* of E . We set

$$\phi^+(E) = \phi_1, \quad \phi^-(E) = \phi_k, \quad m(E) = \sum_i |Z(A_i)| \in \mathbb{R}_{>0}.$$

The real number $m(E)$ is called the *mass* of E . It follows from the definition that the subcategories $\mathcal{P}(\phi) \subset \mathcal{D}$ are abelian categories; the objects of $\mathcal{P}(\phi)$ are said to be *semistable of phase* ϕ , and the simple objects of $\mathcal{P}(\phi)$ are said to be *stable of phase* ϕ . For any interval $I \subset \mathbb{R}$ there is a full subcategory $\mathcal{P}(I) \subset \mathcal{D}$ consisting of objects whose semistable factors have phases in I .

We shall always assume that our stability conditions $\sigma = (Z, \mathcal{P})$ satisfy the *support property* of [23], namely that for some norm $\|\cdot\|$ on $K(\mathcal{D}) \otimes \mathbb{R}$ there is a constant $C > 0$ such that

$$(7.2) \quad \|\gamma\| < C \cdot |Z(\gamma)|$$

for all classes $\gamma \in K(\mathcal{D})$ represented by σ -semistable objects in \mathcal{D} . As explained in [1, Prop. B.4], this is equivalent to assuming that they are *full* and *locally-finite* in the terminology of [2, 4]. We let $\text{Stab}(\mathcal{D})$ denote the set of all such stability conditions on \mathcal{D} .

There is a natural topology on $\text{Stab}(\mathcal{D})$ induced by the metric
(7.3)

$$d(\sigma_1, \sigma_2) = \sup_{0 \neq E \in \mathcal{D}} \left\{ |\phi_{\sigma_2}^-(E) - \phi_{\sigma_1}^-(E)|, |\phi_{\sigma_2}^+(E) - \phi_{\sigma_1}^+(E)|, \left| \log \frac{m_{\sigma_2}(E)}{m_{\sigma_1}(E)} \right| \right\} \in [0, \infty].$$

The following result is proved in [2].

Theorem 7.4. *The space $\text{Stab}(\mathcal{D})$ has the structure of a complex manifold, such that the forgetful map*

$$\pi: \text{Stab}(\mathcal{D}) \longrightarrow \text{Hom}_{\mathbb{Z}}(K(\mathcal{D}), \mathbb{C})$$

taking a stability condition to its central charge, is a local isomorphism.

There are two commuting group actions on $\text{Stab}(\mathcal{D})$ that will be important later. The group of triangulated autoequivalences $\text{Aut}(\mathcal{D})$ acts on $\text{Stab}(\mathcal{D})$ in a rather obvious way: an autoequivalence $\Phi \in \text{Aut}(\mathcal{D})$ acts by

$$\Phi: (Z, \mathcal{P}) \mapsto (Z', \mathcal{P}'), \quad Z'(E) = Z(\Phi^{-1}(E)), \quad \mathcal{P}' = \Phi(\mathcal{P}).$$

There is also an action of the universal cover of the group $\text{GL}^+(2, \mathbf{R})$ of orientation-preserving linear automorphisms of \mathbf{R}^2 . This action does not change the subcategory \mathcal{P} , but acts by post-composition on the central charge, viewed as a map to $\mathbb{C} = \mathbf{R}^2$, with a corresponding adjustment of the grading on \mathcal{P} . This action is not free in general, but there is a subgroup isomorphic to \mathbb{C} which does act freely: an element $t \in \mathbb{C}$ acts by

$$t: (Z, \mathcal{P}) \mapsto (Z', \mathcal{P}'), \quad Z'(E) = e^{-i\pi t} \cdot Z(E), \quad \mathcal{P}'(\phi) = \mathcal{P}(\phi + \text{Re}(t)).$$

Note that for any integer $n \in \mathbb{Z}$, the action of the multiple shift functor $[n]$ coincides with the action of $n \in \mathbb{C}$.

Later we will need the following more precise version of Theorem 7.4.

Proposition 7.5. *Fix a real number $0 < \epsilon \ll 1$. Given a stability condition $\sigma = (Z, \mathcal{P}) \in \text{Stab}(\mathcal{D})$, and a group homomorphism $W: K(\mathcal{D}) \rightarrow \mathbb{C}$ satisfying*

$$|W(E) - Z(E)| < \epsilon \cdot |Z(E)|$$

for all σ -stable objects $E \in \mathcal{D}$, there is a unique stability condition $\sigma' \in \text{Stab}(\mathcal{D})$ with central charge W such that $d(\sigma, \sigma') < \frac{1}{2}$.

Proof. In fact we can take any $0 < \epsilon < \frac{1}{8}$. The support property implies that for any interval $I \subset \mathbf{R}$ of length < 1 , the quasi-abelian categories $\mathcal{P}(I)$ have finite-length. The existence part then follows from the results of [2, Section 7]. The uniqueness is a consequence of [2, Lemma 6.4]. \square

7.6. Walls and chambers. In this section we give some basic results on the wall-and-chamber decomposition of the space of stability conditions. These are well-known, but the proofs in the general setting are not available in the literature. As in the last section, we fix a triangulated category \mathcal{D} and assume that $K(\mathcal{D}) \cong \mathbb{Z}^{\oplus n}$ is free of finite rank.

Proposition 7.6. *Fix an object $E \in \mathcal{D}$. Then*

- (a) *the set of $\sigma \in \text{Stab}(\mathcal{D})$ for which E is σ -stable is open.*
- (b) *the set of $\sigma \in \text{Stab}(\mathcal{D})$ for which E is σ -semistable is closed.*

Proof. For part (a), take a stability condition $\sigma = (Z, \mathcal{P})$ and an object $E \in \mathcal{P}(\phi)$ which is σ -stable. Choose $0 < r \ll 1$ and consider the open ball $B_r(\sigma)$ of radius r centered at σ , with respect to the metric (7.3). By definition of this metric, for $\sigma' = (Z', \mathcal{P}') \in B_r(\sigma)$ there are inclusions

$$\mathcal{P}(\phi) \subset \mathcal{P}'(\phi - r, \phi + r) \subset \mathcal{P}(\phi - 2r, \phi + 2r).$$

Thus E fails to be stable in σ' precisely if there is a triangle $A \rightarrow E \rightarrow B$ whose objects all lie in $\mathcal{P}(\phi - 2r, \phi + 2r)$ and for which $\phi'(A) \geq \phi'(E)$.

The support property implies that for any $M > 0$ there are only finitely many classes $\alpha \in K(\mathcal{D})$ satisfying $|Z(\alpha)| < M$ for which there exist σ -semistable objects of class α . It follows that the set of classes $\alpha \in K(\mathcal{D})$ of objects A as above is finite. Since E is stable in σ we must have $\phi(A) < \phi(E)$ for each such subobject, and so, reducing r if necessary, we can assume that these phase inequalities continue to hold for all $\sigma' \in B_r(\sigma)$. It follows that E is stable for all stability conditions in $B_r(\sigma)$.

Part (b) is immediate: the object E is semistable in σ precisely if $\phi_\sigma^+(E) = \phi_\sigma^-(E)$. By the definition of the metric (7.3) this is a closed condition. \square

Let us now fix a class $\gamma \in K(\mathcal{D})$ and consider stability for objects of this class. Let $\alpha \in K(\mathcal{D})$ be another class which is not proportional to γ . We define

$$W_\gamma(\alpha) \subset \text{Stab}(\mathcal{D})$$

to be the subset of stability conditions $\sigma = (Z, \mathcal{P})$ satisfying the following condition: for some $\phi \in \mathbf{R}$ there is an inclusion $A \subset E$ in the category $\mathcal{P}(\phi)$ such that A and E have classes α and γ respectively. Locally, the subset $W_\gamma(\alpha)$ is contained in the real codimension one submanifold of $\text{Stab}(\mathcal{D})$ defined by the condition $Z(\alpha)/Z(\gamma) \in \mathbf{R}_{>0}$.

Lemma 7.7. *If $B \subset \text{Stab}(\mathcal{D})$ is a compact subset then the set of classes α for which the subset $W_\gamma(\alpha)$ intersects B is finite.*

Proof. The support property for a fixed stability condition σ implies that for any given $M > 0$ there are only finitely many classes $\alpha \in K(\mathcal{D})$ represented by objects of mass $< M$ in σ . On the other hand, the definition of the metric (7.3) shows that the masses of objects of \mathcal{D} vary by a uniformly bounded amount in B , so the same is true if we allow

σ to vary in B . Using compactness again we can assume that M is large enough that $|Z(\gamma)| < M$ for all points in B . But if $\sigma \in W_\gamma(\alpha) \cap B$ then there is an inclusion $A \subset E$ in some $\mathcal{P}(\phi)$, and it follows that A has mass $< M$, and hence has one of finitely-many classes. \square

Consider the complement of the closures

$$\mathcal{C}_\gamma = \text{Stab}(\mathcal{D}) \setminus \bigcup_{\alpha \not\sim \gamma} \bar{W}_\gamma(\alpha)$$

where the union is over classes α which are not proportional to γ . This is the complement of a locally-finite union of closed subsets, hence is open.

We refer to the subsets $W_\gamma(\alpha)$ as *walls* for the class γ , and the connected components of \mathcal{C}_γ will be called *chambers*. The following result shows that the question of whether a given object $E \in \mathcal{D}$ of class γ is stable or semistable has a constant answer for stability conditions in a fixed chamber.

Proposition 7.8. *Let $U \subset \mathcal{C}_\gamma$ be a chamber. If an object E of class γ is (semi)stable for some stability condition $\sigma \in U$ then the same is true for all $\sigma \in U$.*

Proof. We say that an object $E \in \mathcal{D}$ is pseudostable in a stability condition σ if it is semistable and the classes in $K(\mathcal{D})$ of its stable factors are all proportional. The set of points $\sigma \in \text{Stab}(\mathcal{D})$ for which a given object $E \in \mathcal{D}$ is pseudostable is open: indeed, by Lemma 7.6(a), if E is pseudostable for σ then the stable factors of E remain stable in some open neighbourhood of σ , and their phases remain equal since they have proportional classes. The set of points where E is unstable is also open, by Lemma 7.6(b), since it is the complement of the points where E is semistable.

Suppose now that $E \in \mathcal{D}$ has class γ . If E is semistable but not pseudostable then σ must lie on a wall $W_\gamma(\alpha)$. Thus, the subset of points of the chamber U for which E is semistable is both open and closed. Since U is connected, this subset must be either empty or the whole of U .

Assume now that E is semistable for all $\sigma \in U$. As above this implies that E is pseudostable at each $\sigma \in U$. The set of $\sigma \in U$ for which E is stable is then open, by Lemma 7.6(a), and its complement, the set of points for which E is strictly pseudostable is also open, by the argument given above. Hence, if E is stable for some stability condition in U , then it is stable for all of them. \square

7.7. Stability conditions from t-structures. Let \mathcal{D} be a triangulated category. Any stability condition $\sigma = (Z, \mathcal{P})$ on \mathcal{D} has an associated heart

$$\mathcal{A} = \mathcal{P}((0, 1]) \subset \mathcal{D}.$$

It is the extension closure of the subcategories $\mathcal{P}(\phi)$ for $0 < \phi \leq 1$. All nonzero objects of \mathcal{A} are mapped by Z into the semi-closed upper half plane

$$\bar{\mathcal{H}} = \{r \exp(i\pi\phi) : r \in \mathbf{R}_{>0} \text{ and } 0 < \phi \leq 1\} \subset \mathbb{C}.$$

Conversely, given a heart $\mathcal{A} \subset \mathcal{D}$, and a group homomorphism $Z: K(\mathcal{A}) \rightarrow \mathbb{C}$ with this property, then providing some finiteness conditions are satisfied, there is a unique stability condition on \mathcal{D} with heart \mathcal{A} and central charge Z .

In particular, if $\mathcal{A} \subset \mathcal{D}$ is a finite-length heart with n simple objects S_i up to isomorphism, the subset $\text{Stab}(\mathcal{A}) \subset \text{Stab}(\mathcal{D})$ consisting of stability conditions with heart \mathcal{A} is mapped bijectively by π onto the subset

$$\{Z \in \text{Hom}_{\mathbb{Z}}(K(\mathcal{D}), \mathbb{C}) : Z(S_i) \in \bar{\mathcal{H}}\},$$

and is therefore homeomorphic to $\bar{\mathcal{H}}^n$.

The following result shows that tilting controls the way the subsets $\text{Stab}(\mathcal{A})$, for different hearts $\mathcal{A} \subset \mathcal{D}$, are glued together in $\text{Stab}(\mathcal{D})$.

Lemma 7.9. *Let $\mathcal{A} \subset \mathcal{D}$ be a finite-length heart, and suppose that*

$$\sigma = (Z, \mathcal{P}) \in \text{Stab}(\mathcal{D})$$

lies on a unique boundary component of the region $\text{Stab}(\mathcal{A}) \subset \text{Stab}(\mathcal{D})$, so that $\text{Im } Z(S_i) = 0$ for a unique simple object S_i . Assume that the tilted hearts $\mu_{S_i}^{\pm}(\mathcal{A})$ are also finite-length. Then there is a neighbourhood $\sigma \in U \subset \text{Stab}(\mathcal{D})$ such that one of the following holds

- (i) $Z(S_i) \in \mathbf{R}_{<0}$, and $U \subset \text{Stab}(\mathcal{A}) \sqcup \text{Stab}(\mu_{S_i}^+(\mathcal{A}))$,
- (ii) $Z(S_i) \in \mathbf{R}_{>0}$, and $U \subset \text{Stab}(\mathcal{A}) \sqcup \text{Stab}(\mu_{S_i}^-(\mathcal{A}))$.

Proof. This is stated without proof in [3, Lemma 5.5]. A special case is proved in [5, Proposition 2.4], and the general case is proved in exactly the same way. \square

The subsets $\text{Stab}(\mathcal{A}) \subset \text{Stab}(\mathcal{D})$ form a different system of walls and chambers in $\text{Stab}(\mathcal{D})$. The walls consist of points where the subcategory $\mathcal{P}(0)$ contains nonzero objects. To distinguish them from the walls considered in the last subsection they are often referred to as walls of type II. Note that they do not depend on a choice of class $\alpha \in K(\mathcal{D})$.

Suppose that $\mathcal{A} \subset \mathcal{D}$ is a finite-length heart. It follows from Lemma 7.9 that there is a single connected component $\text{Stab}_{\mathcal{A}}(\mathcal{D}) \subset \text{Stab}(\mathcal{D})$ containing all stability conditions whose hearts lie in the connected component $\text{Tilt}_{\mathcal{A}}(\mathcal{D}) \subset \text{Tilt}(\mathcal{D})$. Note however that it is not usually the case that $\text{Stab}_{\mathcal{A}}(\mathcal{D})$ is the union of the chambers $\text{Stab}(\mathcal{B})$ for hearts \mathcal{B} reachable from \mathcal{A} (see [37] for a detailed discussion of this point).

An autoequivalence of \mathcal{D} lying in the subgroup $\text{Aut}_{\mathcal{A}}(\mathcal{D}) \subset \text{Aut}(\mathcal{D})$ of autoequivalences reachable from \mathcal{A} necessarily preserves the connected component $\text{Stab}_{\mathcal{A}}(\mathcal{D})$. The converse is false: the existence of the \mathbb{C} -action shows that the shift functor $[1]$ fixes all connected components of $\text{Stab}(\mathcal{D})$, but it is not generally true that $[1]$ is reachable.

It is easy to see that a triangulated autoequivalence Φ acts trivially on $\text{Stab}_{\mathcal{A}}(\mathcal{D})$ precisely if it fixes a heart $\mathcal{A} \subset \mathcal{D}$ and furthermore fixes pointwise the isomorphism classes of its simple objects. This is equivalent to the condition that Φ acts trivially on the connected component $\text{Tilt}_{\mathcal{A}}(\mathcal{D})$. We say that such autoequivalences are *negligible* with respect to the heart \mathcal{A} .

8. SURFACES AND TRIANGULATIONS

The particular examples of CY_3 categories considered in this paper will be defined using quivers with potential associated to triangulations of marked bordered surfaces. Unfortunately, the non-degenerate ideal triangulations appearing in Section 1.2 will not be sufficient for our purposes. Indeed, to understand the space of stability conditions on our categories, we will need to understand all hearts that are reachable from the standard heart; whilst some of these hearts correspond to non-degenerate ideal triangulations, others correspond to more exotic objects introduced in [10] called tagged triangulations.

8.1. Ideal triangulations. Here we give a brief summary of the relevant definitions concerning triangulations of marked bordered surfaces. The reader can find a more careful treatment in the paper of Fomin, Shapiro and Thurston [10].

A *marked bordered surface* is defined to be a pair (\mathbb{S}, \mathbb{M}) consisting of a compact, connected oriented surface with boundary, and a finite non-empty set $\mathbb{M} \subset \mathbb{S}$ of marked points such that each boundary component of \mathbb{S} contains at least one marked point. Marked points in the interior of \mathbb{S} are called *punctures*; the set of punctures is denoted $\mathbb{P} \subset \mathbb{M}$.

An *arc* in (\mathbb{S}, \mathbb{M}) is a smooth path γ in \mathbb{S} connecting points of \mathbb{M} , whose interior lies in the open subsurface $\mathbb{S} \setminus (\mathbb{M} \cup \partial\mathbb{S})$, and which has no self-intersections in its interior. We moreover insist that γ should not be homotopic, relative to its endpoints, to a single point, or to a path in $\partial\mathbb{S}$ whose interior contains no points of \mathbb{M} . Two arcs are considered to be equivalent if they are related by a homotopy through such arcs.

An *ideal triangulation* of (\mathbb{S}, \mathbb{M}) is defined to be a maximal collection of equivalence classes of arcs for which it is possible to find representatives whose interiors are pairwise disjoint. We refer to the arcs as the *edges* of the triangulation. An example of an ideal triangulation of a disc with 5 marked points on its boundary is depicted in Figure 28; note that it has just two edges. To get something more closely approximating the

intuitive notion of a triangulation of the surface one should add arcs in $\partial\mathbb{S}$ connecting the points of \mathbb{M} .




FIGURE 28. A triangulation of a disc with 5 marked points.

A *face* or *triangle* of an ideal triangulation T is the closure in \mathbb{S} of a connected component of the complement of all arcs of T . A triangle is called *interior* if its intersection with $\partial\mathbb{S}$ is contained in \mathbb{M} . Each interior triangle is topologically a disc, containing either two or three distinct edges of the triangulation.




FIGURE 29. A self-folded triangle.

An interior triangle with just two distinct edges is called a *self-folded triangle*; such a triangle has a *self-folded edge* and an *encircling edge*. The *valency* of a puncture $p \in \mathbb{P}$ with respect to a triangulation T is the number of half-edges of T that are incident with it; a puncture has valency 1 precisely if it is contained in a self-folded triangle.

For various technical reasons, when dealing with triangulations of marked bordered surfaces (\mathbb{S}, \mathbb{M}) , we shall always make the following

Assumption 8.1. *We assume that (\mathbb{S}, \mathbb{M}) is not one of the following surfaces*

- (a) *a sphere with ≤ 5 marked points;*
- (b) *an unpunctured disc with ≤ 3 marked points on the boundary;*
- (c) *a disc with a single puncture and one marked point on its boundary.*

In the cases of a sphere with ≤ 2 punctures, or an unpunctured disc with ≤ 2 marked points, there are no ideal triangulations, and so the theory described below is vacuous. In the cases of an unpunctured disc with 3 marked points, and the surface of case (c), there is a unique ideal triangulation and the theory is trivial and rather degenerate; see Examples 12.1 and 12.3.

The case of a three-punctured sphere is special in that there is an ideal triangulation consisting of two self-folded triangles meeting along a common edge; this plays havoc

with the definition of a tagged triangulation below and for this reason it is better to deal with this case directly: see Section 12.4. Finally, the cases of spheres with 4 or 5 punctures are definitely interesting, but we have to exclude them because the crucial results of Section 9.1 have not been established for these surfaces.

A marked bordered surface (\mathbb{S}, \mathbb{M}) is determined up to diffeomorphism by its genus g , the number of punctures p , and a collection of integers $k_i \geq 1$ encoding the number of marked points on each boundary component. Any ideal triangulation of such a surface has the same number of edges, namely

$$n = 6g - 6 + 3p + \sum_i (k_i + 3).$$

8.2. Flips and pops. Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface satisfying Assumption 8.1. A *signed triangulation* of (\mathbb{S}, \mathbb{M}) is a pair (T, ϵ) consisting of an ideal triangulation T , and a function

$$\epsilon: \mathbb{P} \rightarrow \{\pm 1\}.$$

By a *pop* of a signed triangulation (T, ϵ) we mean the operation of changing the sign $\epsilon(p)$ associated to a puncture $p \in \mathbb{P}$ of valency one. Note that any such puncture lies in the centre of a self-folded triangle of T .

The popping operation generates an equivalence relation on signed triangulations, in which two signed triangulations (T_i, ϵ_i) are equivalent precisely if the underlying triangulations T_i are the same, and the signings ϵ_i differ only at punctures $p \in \mathbb{P}$ of valency one.

It turns out that the equivalence classes for this relation can be explicitly represented by a combinatorial gadget called a *tagged triangulation*. We will explain this in Section 8.3 below, but for now we simply define a tagged triangulation to be an equivalence class of signed triangulations.

Let us introduce notation $\text{Tri}(\mathbb{S}, \mathbb{M})$, $\text{Tri}_\pm(\mathbb{S}, \mathbb{M})$, $\text{Tri}_\bowtie(\mathbb{S}, \mathbb{M})$ for the sets of ideal, signed and tagged triangulations of (\mathbb{S}, \mathbb{M}) respectively. There is a diagram of maps

$$(8.1) \quad \begin{array}{ccc} & & \text{Tri}_\pm(\mathbb{S}, \mathbb{M}) \\ & \nearrow j & \downarrow q \\ \text{Tri}(\mathbb{S}, \mathbb{M}) & \xrightarrow{i} & \text{Tri}_\bowtie(\mathbb{S}, \mathbb{M}) \end{array}$$

where q is the obvious quotient map, and the arrows i and j are embeddings obtained by considering an ideal triangulation as a signed, and hence a tagged triangulation, using the signing $\epsilon \equiv +1$.

Two ideal triangulations T_1 and T_2 are related by a *flip* if they are distinct, and there are edges $e_i \in T_i$ such that $T_1 \setminus \{e_1\} = T_2 \setminus \{e_2\}$. Note that the edges e_1 and e_2 are necessarily

non-self-folded. Conversely, if e is a non-self-folded edge of an ideal triangulation T , it is contained in exactly two triangles of T , and there is a unique ideal triangulation which is the flip of T along e . The flipping operation extends to signed triangulations




FIGURE 30. Flip of a triangulation.

in the obvious way: we flip the underlying triangulation, keeping the signs constant. We say that two tagged triangulations are related by a flip if they can be represented by signed triangulations which differ by a flip.

The sets appearing in the diagram (8.1) can be considered as graphs, with two (ideal, signed, tagged) triangulations being connected by an edge if they differ by a flip. The maps in the diagram then become maps of graphs. The important point is that, unlike the graph $\text{Tri}(\mathbb{S}, \mathbb{M})$ of ideal triangulations, the graph $\text{Tri}_{\bowtie}(\mathbb{S}, \mathbb{M})$ of tagged triangulations is n -regular.

The basic explanation for this regularity is as follows. When a triangulation T_1 contains a self-folded triangle Δ , we cannot flip the self-folded edge f of Δ , so the number of flips that can be performed on T_1 is less than the total number of edges n . On the other hand, if we choose a signing $\epsilon: \mathbb{P} \rightarrow \{\pm 1\}$, and consider the signed triangulation (T_1, ϵ_1) up to the above equivalence relation, then when we flip the encircling edge e of Δ , the puncture p contained in Δ has valency 2 in the new triangulation T_2 , and so there are two inequivalent possible choices for the sign $\epsilon_2(p)$.

It is well-known that any two ideal triangulations of (\mathbb{S}, \mathbb{M}) are related by a finite chain of flips; thus the graph $\text{Tri}(\mathbb{S}, \mathbb{M})$ is always connected [10, Prop. 3.8]. The graph $\text{Tri}_{\bowtie}(\mathbb{S}, \mathbb{M})$ is also connected, except for the case when (\mathbb{S}, \mathbb{M}) is a closed surface with a single puncture $p \in \mathbb{P}$: in that case $\text{Tri}_{\bowtie}(\mathbb{S}, \mathbb{M})$ has two connected components corresponding to the two possible choice of signs $\epsilon(p)$ [10, Prop. 7.10].

8.3. Tagged triangulations. We now explain why the set $\text{Tri}_{\bowtie}(\mathbb{S}, \mathbb{M})$ of tagged triangulations we defined above coincides with the standard version as defined by Fomin, Shapiro and Thurston [10]. This material will not be used in the rest of the paper, and is only logically necessary to justify the above assertions that the graph $\text{Tri}_{\bowtie}(\mathbb{S}, \mathbb{M})$ is connected and n -regular.

Let (\mathbb{S}, \mathbb{M}) be a marked, bordered surface (\mathbb{S}, \mathbb{M}) satisfying Assumption 8.1. A *tagged arc* in (\mathbb{S}, \mathbb{M}) is an arc as defined above, each end of which has been labelled by one

of two labels: plain or tagged. Fix a function $\epsilon: \mathbb{P} \rightarrow \{\pm 1\}$. Given an ordinary arc e , there is a corresponding tagged arc $t_\epsilon(e)$ defined by the following rule:

- (a) If e is not a loop enclosing a once-punctured disc, the underlying arc of $t_\epsilon(e)$ is just e , and an end of e is labelled tagged precisely if it lies at a puncture $p \in \mathbb{P}$ with $\epsilon(p) = -1$.
- (b) If e is a loop based at $m \in \mathbb{M}$, enclosing a disc which contains a single puncture $p \in \mathbb{P}$, then the underlying arc of $t_\epsilon(e)$ is the arc connecting p to m inside the disc. We label the edge adjacent to m tagged precisely if m is a puncture with $\epsilon(m) = -1$, and the edge adjacent to p tagged precisely if $\epsilon(p) = +1$.

By [10, Lemma 9.3], a tagged triangulation in the standard sense considered there is precisely a set of tagged arcs of the form $t_\epsilon(T)$ for some signed triangulation (T, ϵ) . The following result shows that these tagged triangulations are in bijection with the equivalence classes of signed triangulations considered in the last section.

Lemma 8.2. *Suppose that (T_1, ϵ_1) and (T_2, ϵ_2) are signed triangulations. Then $t_{\epsilon_1}(T_1) = t_{\epsilon_2}(T_2)$ if and only if $T_1 = T_2$, and the signings ϵ_1, ϵ_2 differ only at punctures of valency one.*

Proof. Any vertex of valency one lies in the interior of a self-folded triangle, and it is clear from the definition of $t_\epsilon(T)$ that the resulting collection of tagged arcs does not distinguish between the two choices of sign at such a vertex (the enclosing and folded edge of the self-folded triangle get mapped to two taggings of the same arc; changing the sign just exchanges these two).

The converse follows easily from the following observations. Suppose that $\eta \in t_\epsilon(T)$ is a tagged arc, with underlying arc f . Consider tagged arcs $\zeta \in t_\epsilon(T)$ which have the same underlying arc f . If there is no such ζ then we must have $\eta = t_\epsilon(f)$, and the arc f is not contained in a self-folded triangle of T . If there is such a ζ , then there is a self-folded triangle Δ in T with self-folded edge f , and encircling edge e , such that $\{\eta, \zeta\} = \{t_\epsilon(e), t_\epsilon(f)\}$. Moreover, the encircling edge e is completely determined, because the puncture inside Δ is the one at which η and ζ have different markings. \square

We should also check that our definition of when two tagged triangulations differ by a flip coincides with the standard one. Namely, in [10], two tagged triangulations τ_1 and τ_2 are said to be related by a flip if they are distinct, and there are tagged arcs $\eta_i \in \tau_i$ such that $\tau_1 \setminus \{\eta_1\} = \tau_2 \setminus \{\eta_2\}$.

Lemma 8.3. *Two tagged triangulations differ by a flip in the above sense, precisely if they can be represented by signed triangulations differing by a flip.*

Proof. One implication is clear, since if two signed triangulations differ by a flip then by Lemma 8.2, so do the associated tagged triangulations. For the converse, suppose that $\tau_1 \setminus \{\eta_1\} = \tau_2 \setminus \{\eta_2\}$. Then we can write $\tau_1 = t_\epsilon(T_1)$ for some signed triangulation (T_1, ϵ) in such a way that $\eta_1 = t_\epsilon(e_1)$ for some non-self-folded edge e_1 . Flipping this edge gives a different signed triangulation T_2 satisfying $T_1 \setminus \{e_1\} = T_2 \setminus \{e_2\}$. It follows that $\tau_1 \setminus \{\eta_1\} = t_\epsilon(T_2) \setminus t_\epsilon(e_2)$. But the flip of a tagged triangulation in a tagged arc is unique [10, Theorem 7.9]. Thus we have $\tau_2 = t_\epsilon(T_2)$ and $\eta_2 = t_\epsilon(e_2)$. \square

8.4. Edge lattice and quiver. Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface satisfying Assumption 8.1. The *edge lattice* of an ideal triangulation T of (\mathbb{S}, \mathbb{M}) is defined to be the free abelian group $\Gamma(T)$ on the edges of T . We denote by $[e]$ the basis element corresponding to the edge $e \in T$; thus

$$\Gamma(T) = \bigoplus_{e \in T} \mathbb{Z} \cdot [e].$$

For distinct edges $e, f \in T$, we define $c(e, f)$ to be the number of internal triangles of T in which e and f appear as adjacent edges in clockwise order. There is a skew-symmetric form

$$\langle -, - \rangle : \Gamma(T) \times \Gamma(T) \rightarrow \mathbb{Z}$$

given by the formula

$$\langle [e], [f] \rangle = c(f, e) - c(e, f).$$

Note that if e and f are the encircling and self-folded edges of a self-folded triangle then

$$(8.2) \quad c(e, f) = 1 = c(f, e), \quad \langle [e], [f] \rangle = 0.$$

It will be convenient for later purposes to define $c(e, e) = -2$ for all edges e , although of course this has no effect on the form $\langle -, - \rangle$.

We will also need a modified basis $\{e\}$ for the group $\Gamma(T)$ defined as follows:

- (a) if $e \in T$ is not an edge of a self-folded triangle then $\{e\} = [e]$;
- (b) if e and f are respectively the encircling and self-folded edges of a self-folded triangle, then $\{e\} = [e]$ and $\{f\} = [e] + [f]$.

We will see some more intrinsic interpretations of the edge lattice $\Gamma(T)$ later: as a Grothendieck group with its Euler form (Lemma 9.10), and as a homology group with an intersection form (Lemma 10.3). We will also give some explanation for the strange-looking definition of the basis $\{e\}$ (Section 10.2).

Define a map $\kappa : T \rightarrow T$ by setting $\kappa(f) = f$ unless f is a self-folded edge of a self-folded triangle, in which case $\kappa(f) = e$ is the encircling edge of the same triangle. For distinct edges e and f define

$$n(e, f) = \max(0, \langle [\kappa(f)], [e] \rangle) \geq 0.$$

Note that there is a relation

$$(8.3) \quad \langle \{e\}, \{f\} \rangle = \langle [\kappa(e)], [\kappa(f)] \rangle = n(f, e) - n(e, f)$$

for all $e, f \in T$. This is easily checked by noting that when f is a self-folded edge, the basis element $[f]$ lies in the kernel of the form $\langle -, - \rangle$.

To any ideal triangulation T we can now associate a quiver $Q(T)$ whose vertices are the edges of T , and with $n(e, f)$ arrows from vertex e to vertex f . By its definition it has no loops or 2-cycles. In the case of a non-degenerate ideal triangulation of a closed surface it reduces to the quiver considered in Section 1.2.

Remarks 8.4. (a) *If T has a self-folded triangle with edges e and f , then since $\kappa(e) = \kappa(f)$, there is an involution of the quiver $Q(T)$ exchanging the vertices corresponding to these two edges.*

(b) *If e and f are distinct non-self-folded edges, then it is easily checked that $c(e, f)$ and $c(f, e)$ are both nonzero precisely if e and f meet at a puncture of valency 2. Thus if a pair of edges e, f are such that $\kappa(e)$ and $\kappa(f)$ are distinct, and do not meet at a vertex of valency 2, then $n(e, f) = c(\kappa(e), \kappa(f))$.*

8.5. Ordered versions. Sometimes in what follows it will be clearer to work with ordered versions of our basic combinatorial objects: triangulations, quivers, t-structures etc. In this section we gather the necessary definitions; these mostly proceed along the obvious lines.

An ordered ideal triangulation is an ideal triangulation equipped with an ordering of its edges. Similarly, one can consider ordered signed triangulations. By a pop of an ordered signed triangulation we mean the operation which changes the sign $\epsilon(p)$ associated to a puncture $p \in \mathbb{P}$ of valency 1, and which also changes the ordering of the triangulation by transposing the two edges of the self-folded triangle containing p . Two ordered signed triangulations are considered equivalent if they differ by a finite sequence of such pops. By an ordered tagged triangulation we mean an equivalence class of ordered signed tagged triangulations. The map t_ϵ of Section 8.3 respects this equivalence relation, and it follows that we can realise ordered tagged triangulations as ordered collections of tagged arcs.

Two ordered triangulations are related by a flip if the underlying triangulations are related by a flip, and if the orderings of their edges are compatible with the obvious bijection between the edges of the two triangulations. Similarly, one can consider flips of ordered signed triangulations. Two ordered tagged triangulations are related by a flip if they can be represented by ordered signed triangulations that are related by a flip.

An ordered quiver is a quiver equipped with a fixed ordering of its vertices. An ordered triangulation T has an associated ordered quiver $Q(T)$. Remark 8.4(a) shows that the

ordered quiver associated to an ordered triangulation is invariant under transposing the order of the two edges of a self-folded triangle, and it follows that every ordered tagged triangulation also has an associated ordered quiver.

Finally, suppose that \mathcal{D} is a CY_3 triangulated category. By an ordering of a finite-length heart $\mathcal{A} \subset \mathcal{D}$ we mean an ordering of the simple objects of \mathcal{A} . The associated quiver $Q(\mathcal{A})$ is then also ordered in the obvious way. As we explained in the proof of Proposition 7.1, for any simple object $S \in \mathcal{A}$, there is a canonical bijection between the simple objects of the heart \mathcal{A} and those of the tilted heart $\mu_S^\pm(\mathcal{A})$. We say that two ordered hearts $\mathcal{A}, \mathcal{B} \subset \mathcal{D}$ are related by a tilt in a simple object, if the hearts \mathcal{A}, \mathcal{B} are related by such a tilt, and if the orderings on \mathcal{A}, \mathcal{B} are compatible with this canonical bijection.

We denote the graphs of ordered ideal, signed, tagged triangulations by

$$\text{Tri}^\diamond(\mathbb{S}, \mathbb{M}), \quad \text{Tri}_\pm^\diamond(\mathbb{S}, \mathbb{M}), \quad \text{Tri}_{\bowtie}^\diamond(\mathbb{S}, \mathbb{M})$$

respectively. The maps in (8.1) induce maps of the ordered versions in the obvious way. Similarly, given a non-degenerate heart $\mathcal{A} \subset \mathcal{D}$ we use the notation

$$\text{Tilt}_\mathcal{A}^\diamond(\mathcal{D}), \quad \text{Exch}_\mathcal{A}^\diamond(\mathcal{D})$$

for the graphs of ordered reachable finite-length hearts, and the quotient of this graph by the group $\text{Sph}_\mathcal{A}(\mathcal{D})$. We note that these graphs will not be connected in general.

8.6. Mapping class group. By a diffeomorphism of a marked bordered surface (\mathbb{S}, \mathbb{M}) we mean a diffeomorphism of \mathbb{S} which fixes the subset \mathbb{M} , although possibly permuting its elements. The *mapping class group* $\text{MCG}(\mathbb{S}, \mathbb{M})$ is the group of all orientation-preserving diffeomorphisms of (\mathbb{S}, \mathbb{M}) modulo those which are homotopic to the identity through diffeomorphisms of (\mathbb{S}, \mathbb{M}) .

The mapping class group clearly acts on the graphs of (ideal, signed, tagged) triangulations of the surface (\mathbb{S}, \mathbb{M}) , since the edges of such triangulations consist of homotopy classes of arcs.

Proposition 8.5. *Suppose that (\mathbb{S}, \mathbb{M}) is a marked bordered surface which satisfies Assumption 8.1 and which is not one of the following 3 surfaces*

- (a) *a once-punctured disc with 2 or 4 marked points on the boundary;*
- (b) *a twice-punctured disc with 2 marked points on the boundary.*

Then two ideal triangulations of (\mathbb{S}, \mathbb{M}) differ by an element of $\text{MCG}(\mathbb{S}, \mathbb{M})$ precisely if the associated quivers are isomorphic.

Proof. One implication is clear: if two triangulations differ by an orientation-preserving diffeomorphism then they have the same combinatorics and hence the same associated quivers.

For the converse, suppose that two surfaces (S_i, M_i) have ideal triangulations T_i . In [10, Section 13] it is explained how to decompose the quivers $Q(T_i)$ into certain blocks. In the proof of [10, Proposition 14.1] it is shown that if each quiver $Q(T_i)$ has a unique block decomposition, then the quivers $Q(T_i)$ are isomorphic precisely if there is an orientation-preserving diffeomorphism between the surfaces (S_i, M_i) taking one triangulation T_i to the other. Weiwen Gu [16] has classified all quivers which have more than one block decomposition. The only examples corresponding to combinatorially-distinct triangulations of the same surface occur when (S, M) is one of the three cases listed in the statement of the Proposition, or a sphere with 3 or 4 punctures. These last two cases are already excluded by Assumption 8.1. \square

The mapping class group of (S, M) acts on the set $\mathbb{P} \subset M$ of punctures in the obvious way, and we define the *signed mapping class group* to be the corresponding semi-direct product

$$(8.4) \quad \text{MCG}^\pm(S, M) = \text{MCG}(S, M) \ltimes \mathbb{Z}_2^{\mathbb{P}}.$$

This group acts on the set of signed triangulations, with the $\mathbb{Z}_2^{\mathbb{P}}$ part acting by changing the signs $\epsilon(p) \in \{\pm 1\}$ associated to the punctures. This action clearly descends to an action on tagged triangulations. It is an immediate consequence of Proposition 8.5 that the quivers associated to two signed or tagged triangulations of (S, M) are isomorphic precisely if they differ by an element of the signed mapping class group.

8.7. Free action on ordered triangulations. One reason to introduce ordered triangulations is the following result.

Proposition 8.6. *Suppose that (S, M) satisfies Assumption 8.1 and is not one of the following 3 surfaces:*

- (i) *an unpunctured disc with 4 points on its boundary;*
- (ii) *an annulus with one marked point on each boundary component;*
- (iii) *a closed torus with a single puncture.*

Then the action of the mapping class group $\text{MCG}(S, M)$ on the set $\text{Tri}^\diamond(S, M)$ of ordered ideal triangulations is free. Similarly, the actions of the signed mapping class group $\text{MCG}^\pm(S, M)$ on the sets $\text{Tri}_\pm^\diamond(S, M)$ and $\text{Tri}_\infty^\diamond(S, M)$ are free.

Proof. For the case of ideal triangulations, we must show that an orientation-preserving diffeomorphism g of (S, M) which fixes the edges of such a triangulation T is homotopic to the identity, through diffeomorphisms of (S, M) .

Suppose first that g fixes every triangle of T . Since g then induces an orientation-preserving diffeomorphism of each triangle, it follows that g preserves the orientation of each edge of T . Moreover, since every triangle contains at least one edge, we see that g preserves each connected component of $\partial S \setminus M$. Now $\text{Diff}(I, \partial I)$ is contractible, and the

edges are disjoint in their interiors, so we can isotope g so that it fixes all edges of T , and all components of $\partial\mathbb{S}$ pointwise. The result then follows from the fact [31, Theorem B] that the group $\text{Diff}(D^2, \partial D^2)$ is contractible.

Suppose instead that there is some triangle Δ such that $g(\Delta) \neq \Delta$. Then the edges of Δ coincide with those of $g^i(\Delta)$ for all $i \in \mathbb{Z}$. Since any edge occurs in the boundary of at most 2 triangles, it follows that $g^2(\Delta) = \Delta$. Moreover, the surface \mathbb{S} is completely covered by the two triangles Δ and $g(\Delta)$, since passing through an edge of Δ takes us into $g(\Delta)$ and vice versa. We then obtain the three possibilities listed, according to whether the closures of the triangles Δ and $g(\Delta)$ meet in 1, 2 or 3 edges.

The extension to signed triangulations is obvious. For the case of tagged triangulations, suppose that an element of $g \in \text{MCG}^\pm(\mathbb{S}, \mathbb{M})$ fixes an ordered tagged triangulation τ , which we view as an equivalence-class of signed triangulations (T, ϵ) . Note that the action of g on $\text{Tri}_\pm^\circ(\mathbb{S}, \mathbb{M})$ commutes with the flipping operation. Thus we can reduce to the case when T has no self-folded triangles. Then g must fix the signed triangulation (T, ϵ) , since this is the only signed triangulation in the equivalence-class τ . Hence g is the identity. \square

We note that the excluded cases in Proposition 8.6 are essentially the same as those in Proposition 6.6 (the case of an unpunctured disc with 3 marked points would appear above were it not already excluded by Assumption 8.1). This is not a coincidence: generic automorphisms of the space $\text{Quad}(\mathbb{S}, \mathbb{M})$ correspond to automorphisms of (\mathbb{S}, \mathbb{M}) which preserve a horizontal strip decomposition together with an ordering of the horizontal strips. Any such automorphism will also preserve an ordered version of the WKB triangulation of Section 10.1.

9. THE CATEGORY ASSOCIATED TO A SURFACE

In this section we introduce the particular examples of CY_3 categories that appear in our main Theorems. They are indexed by diffeomorphism classes of marked bordered surfaces (\mathbb{S}, \mathbb{M}) . We also provide the combinatorial underpinning of our main theorems, by giving a precise correspondence between tagged triangulations of (\mathbb{S}, \mathbb{M}) and t-structures on the corresponding category $\mathcal{D}(\mathbb{S}, \mathbb{M})$. Throughout this section we rely heavily on the work of D. Labardini-Fragoso.

9.1. Some results of Labardini-Fragoso. Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface satisfying Assumption 8.1. Let T be an ideal triangulation of (\mathbb{S}, \mathbb{M}) . Labardini-Fragoso [24] defined a reduced potential $W(T)$ on the quiver $Q(T)$ introduced in Section 8.4, depending also on some nonzero scalar constants $x_p \in k \setminus \{0\}$, one for each puncture $p \in \mathbb{P}$. We shall always take these scalars to be defined by a signing; thus we take

a signed triangulation (T, ϵ) and consider the quiver with potential $(Q(T), W(T, \epsilon))$ obtained by setting $x_p = \epsilon(p)$.

In the case when T is non-degenerate, the resulting potential reads

$$(9.1) \quad W(T, \epsilon) = \sum_f T(f) - \sum_p \epsilon(p) C(p),$$

where $T(f)$ and $C(p)$ are the cycles in $Q(T)$ defined in Section 1.2. In the presence of punctures of valency ≤ 2 the recipe becomes more complicated. The explicit form of the potential will not be important for what follows, and we refer to [24, Section 3] for details. What will be important are the invariance properties under flips and pops which we now describe.

The case of flips is dealt with by the following result of Labardini-Fragoso [24, Theorem 30]. To be absolutely clear we state it for ordered triangulations.

Theorem 9.1. *Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface satisfying Assumption 8.1. Suppose that two ordered signed triangulations (T_i, ϵ_i) of (\mathbb{S}, \mathbb{M}) are related by a flip in a non-self-folded edge e . Then, up to right-equivalence, the ordered quivers with potential $(Q(T_i), W(T_i, \epsilon_i))$ are related by a mutation at the corresponding vertex.*

We now move on to the case of pops. As we remarked in Section 8.5, the popping symmetry of Remark 8.4(a) implies that an ordered tagged triangulation has an associated ordered quiver. The following result [7, Theorem 6.1] extends this statement to quivers with potential.

Theorem 9.2. *Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface satisfying Assumption 8.1. Suppose that two ordered signed triangulations (T, ϵ_i) of (\mathbb{S}, \mathbb{M}) are related by a pop. Then the associated ordered quivers with potential $(Q(T), W(T, \epsilon_i))$ are right-equivalent.*

We conclude that every ordered tagged triangulation has an associated ordered quiver with potential, well-defined up to right-equivalence. If we think of a tagged triangulation τ as a collection of tagged arcs as in Section 8.3, then we can say that there is an associated quiver with potential $(Q(\tau), W(\tau))$, well-defined up to right-equivalence, whose vertices are in natural bijection with these tagged arcs.

To avoid all technical difficulties we shall mostly work with the following class of surfaces. We return to some of the exceptional cases in Section 11.6.

Definition 9.3. *We say that a marked bordered surface (\mathbb{S}, \mathbb{M}) is amenable if*

- (a) (\mathbb{S}, \mathbb{M}) satisfies Assumption 8.1;
- (b) (\mathbb{S}, \mathbb{M}) is not one of the 3 surfaces listed in Proposition 8.5;
- (c) (\mathbb{S}, \mathbb{M}) is not one of the 3 surfaces listed in Proposition 8.6;
- (d) (\mathbb{S}, \mathbb{M}) is not a closed surface with a single puncture.

Note for example that (\mathbb{S}, \mathbb{M}) is amenable if $g(\mathbb{S}) > 0$ and $|\mathbb{M}| > 1$.

Recall from Section 7.4 the definition of a non-degenerate quiver with potential.

Theorem 9.4. *Suppose that the marked bordered surface (\mathbb{S}, \mathbb{M}) is amenable. Then the quiver with potential $(Q(T), W(T, \epsilon))$ associated to any signed triangulation of (\mathbb{S}, \mathbb{M}) is non-degenerate.*

Proof. Up to right-equivalence the quiver with potential $(Q(T), W(T, \epsilon))$ associated to a signed triangulation depends only on the corresponding tagged triangulation. Since tagged triangulations can be flipped in any edge, and such flips can be expressed as flips of signed triangulations to which Theorem 9.1 applies, we conclude that every mutation of $(Q(T), W(T, \epsilon))$ is of the same type. The result follows, since by definition, none of the quivers $Q(T)$ has loops or oriented 2-cycles. \square

9.2. Definition of the category. Let (\mathbb{S}, \mathbb{M}) be an amenable marked bordered surface. In this section we introduce the associated CY_3 triangulated category $\mathcal{D}(\mathbb{S}, \mathbb{M})$, well-defined up to k -linear triangulated equivalence.

Given a signed triangulation (T, ϵ) of (\mathbb{S}, \mathbb{M}) , let us write

$$\mathcal{D}(T, \epsilon) = \mathcal{D}(Q(T), W(T, \epsilon))$$

for the CY_3 category defined by the quiver with potential considered in the last subsection. This category comes equipped with a standard heart

$$\mathcal{A}(T, \epsilon) \subset \mathcal{D}(T, \epsilon),$$

whose simple objects S_e are indexed by the edges e of the triangulation T .

Combining Theorem 7.3 and Theorem 9.1 immediately gives

Theorem 9.5. *Suppose that two signed triangulations (T_i, ϵ_i) of (\mathbb{S}, \mathbb{M}) differ by a flip in an edge e . Then there is a canonical pair of k -linear triangulated equivalences*

$$\Phi_{\pm}: \mathcal{D}(T_1, \epsilon_1) \rightarrow \mathcal{D}(T_2, \epsilon_2),$$

satisfying $\Phi_{\pm}(\mathcal{A}(T_1, \epsilon_1)) = \mu_{S_e}^{\pm}(\mathcal{A}(T_2, \epsilon_2))$ and inducing the natural bijection on simple objects.

Similarly, Theorem 9.2 implies

Theorem 9.6. *Suppose that two signed triangulations (T, ϵ_i) of (\mathbb{S}, \mathbb{M}) are related by a pop at a puncture $p \in \mathbb{P}$. Then there is a k -linear triangulated equivalence*

$$\Psi: \mathcal{D}(T, \epsilon_1) \rightarrow \mathcal{D}(T, \epsilon_2)$$

which identifies the standard hearts, and exchanges the two simple objects S_e and S_f corresponding to the two edges of the self-folded triangle containing p .

Since the graph $\text{Tri}_\infty(\mathbb{S}, \mathbb{M})$ is connected, these two results show that, up to k -linear triangulated equivalence, the category $\mathcal{D}(T, \epsilon)$ depends only on the surface (\mathbb{S}, \mathbb{M}) and not on the chosen signed triangulation. Thus we can associate to the surface (\mathbb{S}, \mathbb{M}) a CY_3 triangulated category

$$\mathcal{D} = \mathcal{D}(\mathbb{S}, \mathbb{M}).$$

In fact it will be important to make a slightly stronger statement, as we now explain. Each category $\mathcal{D}(T, \epsilon)$ comes with a distinguished connected component of its tilting graph, namely the one containing the standard heart. Moreover, if (T_i, ϵ_i) are two signed triangulations of (\mathbb{S}, \mathbb{M}) then by composing the equivalences of Theorems 9.5 and 9.6 we obtain equivalences $\mathcal{D}(T_1, \epsilon_1) \cong \mathcal{D}(T_2, \epsilon_2)$ which identify these connected components. Thus the category \mathcal{D} comes equipped with a distinguished connected component

$$\text{Tilt}_\Delta(\mathcal{D}) \subset \text{Tilt}(\mathcal{D}).$$

Adapting the general notation from Section 7, we write

$$\text{Aut}_\Delta(\mathcal{D}) \subset \text{Aut}(\mathcal{D})$$

for the group of autoequivalences of \mathcal{D} which preserve this connected component; such autoequivalences will be called *reachable*. We write

$$\text{Nil}_\Delta(\mathcal{D}) \subset \text{Aut}_\Delta(\mathcal{D})$$

for the autoequivalences which act trivially on $\text{Tilt}_\Delta(\mathcal{D})$; we call these autoequivalences *negligible*.

Negligible autoequivalences fix the simple objects of all the hearts in our distinguished component $\text{Tilt}_\Delta(\mathcal{D})$. They will also act trivially on the corresponding distinguished connected component $\text{Stab}_\Delta(\mathcal{D})$. For this reason, it is useful to consider the quotient group

$$\mathcal{A}ut_\Delta(\mathcal{D}) = \text{Aut}_\Delta(\mathcal{D}) / \text{Nil}_\Delta(\mathcal{D})$$

which acts effectively on these spaces.

Remark 9.7. *We could instead consider defining a category $\mathcal{D}^\dagger(\mathbb{S}, \mathbb{M})$ by using uncompleted Ginzburg algebras, rather than the complete ones we are using here. It seems likely that if (\mathbb{S}, \mathbb{M}) is amenable the resulting category would be equivalent to $\mathcal{D}(\mathbb{S}, \mathbb{M})$. As evidence for this, note that in the case when \mathbb{S} has non-empty boundary the natural map $J^\dagger(Q, W) \rightarrow J(Q, W)$ from the uncompleted Jacobi algebra to the complete version is an isomorphism [7, Theorem 5.7]. When the surface is not amenable this statement can definitely fail. For example, when (\mathbb{S}, \mathbb{M}) is a closed torus with a single puncture the algebra $J^\dagger(Q, W)$ is infinite-dimensional, whereas $J(Q, W)$ is finite-dimensional [25, Example 8.2]. See also Example 12.4 below for the case of the three-punctured sphere.*

9.3. T-structures and autoequivalences. Let (\mathbb{S}, \mathbb{M}) be an amenable marked bordered surface, and let $\mathcal{D} = \mathcal{D}(\mathbb{S}, \mathbb{M})$ be the associated CY_3 triangulated category. In this section we study the distinguished connected component $\text{Tilt}_\Delta(\mathcal{D})$ of the tilting graph of \mathcal{D} , and the corresponding group $\text{Aut}_\Delta(\mathcal{D})$ of reachable autoequivalences.

Recall from Section 7.3 that there is a subgroup

$$\mathrm{Sph}_\Delta(\mathcal{D}) \subset \mathrm{Aut}_\Delta(\mathcal{D}),$$

generated by the twist functors Tw_{S_i} in the simple objects of any heart $\mathcal{A} \in \mathrm{Tilt}_\Delta(\mathcal{D})$. We write

$$\mathcal{Sph}_\Delta(\mathcal{D}) \subset \mathcal{Aut}_\Delta(\mathcal{D}),$$

for the corresponding subgroup of $\mathcal{Aut}_\Delta(\mathcal{D})$. The group $\mathrm{Sph}_\Delta(\mathcal{D})$ acts on the tilting graph $\mathrm{Tilt}(\mathcal{D})$, and Proposition 7.1 implies that this action preserves the connected component $\mathrm{Tilt}_\Delta(\mathcal{D})$. We call the quotient graph

$$\mathrm{Exch}_\Delta(\mathcal{D}) = \mathrm{Tilt}_\Delta(\mathcal{D}) / \mathrm{Sph}_\Delta(\mathcal{D})$$

the *heart exchange graph* of (\mathbb{S}, \mathbb{M}) . There is also an ordered version $\mathrm{Exch}_\Delta^\diamond(\mathcal{D})$ defined in the obvious way.

The following result gives the basic link between triangulations of the surface (\mathbb{S}, \mathbb{M}) and t-structures in the corresponding category $\mathcal{D}(\mathbb{S}, \mathbb{M})$.

Theorem 9.8. *There are isomorphisms of graphs*

$$\mathrm{Tri}_\bowtie(\mathbb{S}, \mathbb{M}) \cong \mathrm{Exch}_\Delta(\mathcal{D}), \quad \mathrm{Tri}_\bowtie^\diamond(\mathbb{S}, \mathbb{M}) \cong \mathrm{Exch}_\Delta^\diamond(\mathcal{D}).$$

Proof. Corollary 3.6 of [6] and Theorem 5.6 of [19] together imply that the heart exchange graph $\mathrm{Exch}_\Delta(\mathcal{D})$ is isomorphic to the cluster exchange graph. On the other hand, the cluster exchange graph was shown to be isomorphic to the tagged triangulation graph $\mathrm{Tri}_\bowtie(\mathbb{S}, \mathbb{M})$ in [10, Theorem 7.11] and [11]. Both of these isomorphisms are constructed in such a way that they lift to maps of the corresponding ordered graphs, and it follows that these maps are also isomorphisms. \square

We note that the isomorphisms of Theorem 9.8 have the property that if a tagged triangulation is represented by a signed triangulation (T, ϵ) , then the corresponding heart $\mathcal{A} \subset \mathcal{D}$ is the image of the standard heart $\mathcal{A}(T, \epsilon)$ under an equivalence

$$\mathcal{D}(T, \epsilon) \cong \mathcal{D}.$$

The obvious generalization to the ordered versions also holds. We can use this result to prove⁹ a result on the structure of the group $\mathcal{Aut}_\Delta(\mathcal{D})$.

Theorem 9.9. *Assume that (\mathbb{S}, \mathbb{M}) is amenable. Then there is a short exact sequence*

$$1 \longrightarrow \mathcal{Sph}_\Delta(\mathcal{D}) \longrightarrow \mathcal{Aut}_\Delta(\mathcal{D}) \longrightarrow \mathrm{MCG}^\pm(\mathbb{S}, \mathbb{M}) \longrightarrow 1.$$

⁹The method of proof was explained to us by Alastair King.

Proof. Consider the n -fold free product group

$$F_n = \mathbb{Z}_2 * \cdots * \mathbb{Z}_2 = \langle \mu_1, \dots, \mu_n : \mu_i^2 = 1 \rangle.$$

There is an obvious action of the symmetric group Sym_n permuting the generators μ_i , and we also consider the semi-direct product

$$G_n = \text{Sym}_n \ltimes F_n.$$

The set $\text{Tri}_{\bowtie}^{\diamond}(\mathbb{S}, \mathbb{M})$ of ordered, tagged triangulations has a natural action of the group G_n : the generator μ_i acts by flipping the i th edge of an ordered triangulation, and the group Sym_n acts by permuting the ordering of the edges. This action is transitive, because, with our assumptions on \mathbb{S} , the graph of unordered tagged triangulations $\text{Tri}_{\bowtie}(\mathbb{S}, \mathbb{M})$ is connected.

In a similar way, the set $\text{Exch}_{\Delta}^{\diamond}(\mathcal{D})$ of ordered reachable hearts has a transitive action of the group G_n . This time the generator μ_i acts by tilting an ordered heart at the i th simple object. Note that the left and right tilts co-incide on the exchange graph by Proposition 7.1(a), and the relation $\mu_i^2 = 1$ follows from the relation (7.1).

Consider the set \mathcal{Q}_n of right-equivalence classes of ordered quivers with n vertices. It carries an action of the group G_n , where the generator μ_i acts by mutation at the i th vertex. We now have a commutative diagram of G_n -equivariant maps of sets

$$(9.2) \quad \begin{array}{ccc} \text{Tri}_{\bowtie}^{\diamond}(\mathbb{S}, \mathbb{M}) & \xrightarrow{\theta} & \text{Exch}_{\Delta}^{\diamond}(\mathcal{D}) \\ & \searrow p \quad \swarrow q & \\ & \mathcal{Q}_n & \end{array}$$

where p sends an ordered tagged triangulation (T, ϵ) to the associated ordered quiver $Q(T)$, the map q sends an ordered heart $\mathcal{A} \subset \mathcal{D}$ to the associated quiver $Q(\mathcal{A})$, and θ is the bijection of Theorem 9.8.

The signed mapping class group $\text{MCG}^{\pm}(\mathbb{S}, \mathbb{M})$ acts freely on $\text{Tri}_{\bowtie}^{\diamond}(\mathbb{S}, \mathbb{M})$ by Proposition 8.6 and obviously commutes with the G_n -action. Proposition 8.5 shows that the orbits for this action are precisely the fibres of p . It then follows from the transitivity of the G_n -action that $\text{MCG}^{\pm}(\mathbb{S}, \mathbb{M})$ can be identified with the group of automorphisms of the set $\text{Tri}_{\bowtie}^{\diamond}(\mathbb{S}, \mathbb{M})$ which commute with the G_n -action and preserve the map p .

By the definition of a negligible autoequivalence, the group $\mathcal{A}ut_{\Delta}(\mathcal{D})$ acts freely on the graph $\text{Tilt}_{\Delta}^{\diamond}(\mathcal{D})$ of ordered reachable hearts. Quotienting $\mathcal{A}ut_{\Delta}(\mathcal{D})$ by the normal subgroup $\mathcal{S}ph_{\Delta}(\mathcal{D})$, we therefore obtain a free action of the group

$$(9.3) \quad \mathcal{A}ut_{\Delta}(\mathcal{D}) / \mathcal{S}ph_{\Delta}(\mathcal{D})$$

on the set $\text{Exch}_{\Delta}^{\diamond}(\mathcal{D})$, commuting with the G_n -action. To complete the proof we must show that the orbits of $\mathcal{A}ut_{\Delta}(\mathcal{D})$ are precisely the fibres of the map q .

Suppose that two ordered hearts $\mathcal{A}_i \subset \mathcal{D}$ lie in the same fibre of q . Under the bijection θ , these hearts correspond to ordered tagged triangulations (T_i, ϵ_i) lying in the same fibre of p . It follows that they differ by an element $g \in \text{MCG}^\pm(\mathbb{S}, \mathbb{M})$. We claim that the ordered quivers with potential $(Q(T_i), W(T_i, \epsilon_i))$ are right-equivalent up to scale, meaning that there is a nonzero scalar $\lambda \in k^*$ and a right-equivalence

$$(Q(T_1, \cdot), W(T_1, \epsilon_1)) \sim (Q(T_2), \lambda W(T_2, \epsilon_2)).$$

In particular it follows that there are equivalences

$$\mathcal{D}(T_1, \epsilon_1) \cong \mathcal{D}(T_2, \epsilon_2)$$

preserving the standard hearts. The remark following Theorem 9.8 then shows that the hearts \mathcal{A}_i differ by an autoequivalence of \mathcal{D} , and the result then follows.

When $g \in \text{MCG}(\mathbb{S}, \mathbb{M})$ the claim is obvious since the quiver with potential associated to a signed triangulation depends only on the combinatorial structure of the triangulation. For general $g \in \text{MCG}^\pm(\mathbb{S}, \mathbb{M})$ the claim follows from the statement that the quiver with potential $(Q(T), W(T, \epsilon))$ is independent of the signing ϵ , up to scaling and right-equivalence. A proof of this statement will appear in [12] (see [7, Prop. 10.4]). When \mathbb{S} has non-empty boundary, no scaling is necessary, and the claim follows from [7, Prop. 10.2]. \square

9.4. Grothendieck group. Let (\mathbb{S}, \mathbb{M}) be an amenable marked bordered surface. Recall from Section 8.4 the definition of the edge lattice $\Gamma(T)$ associated to an ideal triangulation of (\mathbb{S}, \mathbb{M}) .

Lemma 9.10. *Let (T, ϵ) be a signed triangulation of (\mathbb{S}, \mathbb{M}) . Then there is an isomorphism of abelian groups*

$$\lambda: \Gamma(T) \rightarrow K(\mathcal{D}(T, \epsilon)),$$

such that for each edge e , the basis element $\{e\}$ is mapped to the class of the corresponding simple object S_e . This map takes the form $\langle -, - \rangle$ on $\Gamma(T)$ to the Euler form on $K(\mathcal{D}(T, \epsilon))$.

Proof. For any reduced quiver with potential (Q, W) , the Grothendieck group of the category $\mathcal{D}(Q, W)$ is identified with that of its canonical heart $\mathcal{A} = \mathcal{A}(Q, W)$, and is therefore the free abelian group on the vertices of Q . The CY_3 condition implies that the Euler form is given by skew-symmetrising the adjacency matrix of the quiver. The result is then immediate from (8.3). \square

Suppose that (T_i, ϵ_i) are signed triangulations of (\mathbb{S}, \mathbb{M}) differing by a flip in an edge e . We then use the natural bijection between the edges of T_1 and T_2 to identify these two sets. Theorem 9.5 gives equivalences

$$\Phi_\pm: \mathcal{D}(T_1, \epsilon_1) \cong \mathcal{D}(T_2, \epsilon_2),$$

which induce isomorphisms ϕ_{\pm} on the Grothendieck groups. We have the following explicit formulae for these maps.

Lemma 9.11. *Define maps F_{\pm} by the commutative diagram*

$$\begin{array}{ccc} \Gamma(T_1) & \xrightarrow{F_{\pm}} & \Gamma(T_2) \\ \lambda_1 \downarrow & & \downarrow \lambda_2 \\ K(\mathcal{D}(T_1, \epsilon_1)) & \xrightarrow{\phi_{\pm}} & K(\mathcal{D}(T_2, \epsilon_2)) \end{array}$$

where the λ_i are the maps of Lemma 9.10. Then for all edges f we have

$$(9.4) \quad F_+(\{f\}) = \{f\} + n(e, f) \{e\}, \quad F_-(\{f\}) = \{f\} + n(f, e) \{e\},$$

where $n(-, -)$ is computed in the triangulation T_2 , and we set $n(e, e) = -2$.

Proof. According to Theorem 9.5, the simple objects of the canonical heart $\mathcal{A}(T_1, \epsilon_1)$ are mapped by the functor Φ_{\pm} to the simple objects of the tilted heart $\mu_{S_e}^{\pm}(\mathcal{A}(T_2, \epsilon_2))$. The simple objects of the mutated heart $\mu_{S_e}^+(\mathcal{A}(T_2, \epsilon_2))$ were listed in the proof of Proposition 7.1. The first formula of (9.4) then follows because the extension groups between the simple objects in $\mathcal{A}(T_2, \epsilon_2)$ are based by the arrows in the quiver $Q(T_2)$.

To prove the second formula in (9.4) note that by (8.3) it differs from the first by a reflection in the element $\{e\}$ with respect to the form $\langle -, - \rangle$. By Lemma 9.10 this corresponds under the isometry λ_2 to the action of the twist functor Tw_{S_e} on $K(\mathcal{D}(T_2, \epsilon_2))$. The result therefore follows from Proposition 7.1(a). \square

We have a similar result for pops. Suppose that (T, ϵ_i) are signed triangulations of (\mathbb{S}, \mathbb{M}) differing by a pop at a puncture $p \in \mathbb{P}$. Theorem 9.6 gives an equivalence

$$\Psi: \mathcal{D}(T, \epsilon_1) \cong \mathcal{D}(T, \epsilon_2),$$

which induces an isomorphism ψ on the Grothendieck groups.

Lemma 9.12. *Define a map F by the commutative diagram*

$$\begin{array}{ccc} \Gamma(T) & \xrightarrow{F} & \Gamma(T) \\ \lambda_1 \downarrow & & \downarrow \lambda_2 \\ K(\mathcal{D}(T, \epsilon_1)) & \xrightarrow{\psi} & K(\mathcal{D}(T, \epsilon_2)) \end{array}$$

where the λ_i are the maps of Lemma 9.10. Then the map F exchanges the two elements $\{e\}$ and $\{f\}$ corresponding to the edges of the self-folded triangle containing p , and fixes all other elements of this basis.

Proof. Immediate from Theorem 9.6. \square

9.5. An unpleasant Lemma. Let (\mathbb{S}, \mathbb{M}) be an amenable marked bordered surface. Suppose that (T_i, ϵ) are two signed triangulations related by a flip in an edge e . The purpose of this section is to write the maps F_\pm of Lemma 9.11 in terms of the original basis elements $[e]$ of the lattices $\Gamma(T_i)$. As before, we use the natural bijection between the edges of T_1 and T_2 to identify these two sets.

Lemma 9.13. *The maps F_\pm of Lemma 9.11 satisfy*

$$(9.5) \quad F_+([f]) = [f] + c(e, f) [e], \quad F_-([f]) = [f] + c(f, e) [e],$$

where $c(-, -)$ is computed in the triangulation T_2 , and we set $c(e, e) = -2$.

Proof. As in the proof of Lemma 9.11, it is enough to check the result for F_+ since the two formulae differ by a reflection in the element $[e] = \{e\}$. Recall from Section 8.4 the map κ , which sends an edge to itself, unless it is self-folded, in which case it sends it to the corresponding encircling edge. We have the relations

$$(9.6) \quad \{f\} = [f] + \rho_i(f) [\kappa_i(f)], \quad i = 1, 2,$$

where $\rho_i(f) = 0$ unless edge f is self-folded in the triangulation T_i , in which case it is equal to 1. By definition, the edge e is not self-folded in either triangulation, so $\rho_i(e) = 0$ and $\kappa_i(e) = e$.

We note two basic facts which we will use in the proof. Firstly, if f is self-folded in T_1 then f fails to be self-folded in T_2 if and only if $e = \kappa_1(f)$. Secondly, if f is not self-folded in T_1 then it is self-folded in T_2 precisely if e and f meet in T_1 at a puncture of valency 2.

The formula of the statement certainly defines some isomorphism; we must just check that it agrees with the formula of Lemma 9.11. Substituting (9.6) into (9.5) gives

$$\begin{aligned} F_+(\{f\}) &= [f] + \rho_1(f) [\kappa_1(f)] + (c(e, f) + \rho_1(f) c(e, \kappa_1(f))) [e]. \\ &= \{f\} - \rho_2(f) [\kappa_2(f)] + \rho_1(f) [\kappa_1(f)] + (c(e, f) + \rho_1(f) c(e, \kappa_1(f))) \{e\}. \end{aligned}$$

We now claim that

$$-\rho_2(f) [\kappa_2(f)] + \rho_1(f) [\kappa_1(f)] = (\rho_1(f) \delta_{e, \kappa_1(f)} - \rho_2(f) \delta_{e, \kappa_2(f)}) \{e\}.$$

To see this, note that both sides are zero unless f is self-folded in one of the T_i . If f is self-folded in both, then e is not equal to $\kappa_1(f)$ or $\kappa_2(f)$, and since $\kappa_1(f) = \kappa_2(f)$, both sides are still zero. If f is self-folded in T_1 but not in T_2 then necessarily $e = \kappa_1(f)$, and both sides return $\{e\}$. Similarly, if f is self-folded in T_2 but not in T_1 then $e = \kappa_2(f)$ and both sides return $-\{e\}$.

Thus it remains to show that

$$n(e, f) = c(e, f) + \rho_1(f) c(e, \kappa_1(f)) + \rho_1(f) \delta_{e, \kappa_1(f)} - \rho_2(f) \delta_{e, \kappa_2(f)},$$

where $c(-, -)$ and $n(-, -)$ are always computed in the triangulation T_2 . Write $m(e, f)$ for the expression on the right. If $e = f$ then since e is not self-folded in either triangulation we have $m(e, e) = c(e, e) = n(e, e) = -2$. Thus we assume that $e \neq f$. We

proceed by a case-by-case analysis according to whether f is self-folded in each of the two triangulations T_i .

Case (a). $\rho_1(f) = 0, \rho_2(f) = 0$. In this case we have $n(e, f) = c(e, f) = m(e, f)$.

Case (b). $\rho_1(f) = 0, \rho_2(f) = 1$. Then $e = \kappa_2(f)$ so $n(e, f) = 0$ and $m(e, f) = 1 - 1 = 0$.

Case (c). $\rho_1(f) = 1, \rho_2(f) = 0$. Then e and f must meet at a vertex of valency 2 in T_2 , so $c(e, f) = c(f, e) = 1$ and $n(e, f) = 0$. But then $\kappa_1(f) = e$, so $m(e, f) = 1 - 2 + 1 = 0$.

Case (d). $\rho_1(f) = 1, \rho_2(f) = 1$. Then $\kappa_1(f) = \kappa_2(f)$ and $e \neq \kappa_2(f)$, so using Remark 8.4(b) we have $m(e, f) = c(e, \kappa_2(f)) = n(e, f)$. \square

10. FROM DIFFERENTIALS TO STABILITY CONDITIONS

In this last part of the paper we shall prove our main theorems, by combining the geometry of Sections 2–6 with the algebra and combinatorics of Sections 7–9. In this first section we explain the basic link between quadratic differentials and stability conditions, following the ideas of Gaiotto, Moore and Neitzke [14, Section 6].

Our starting point is the observation that a complete and saddle-free differential $\phi \in \text{Quad}(\mathbb{S}, \mathbb{M})$ determines an ideal triangulation $T(\phi)$ of the surface (\mathbb{S}, \mathbb{M}) up to the action of the mapping class group $\text{MCG}(\mathbb{S}, \mathbb{M})$. We then go on to study how this triangulation changes as we cross between different connected components of the open subset of saddle-free differentials.

10.1. WKB triangulation. Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface. Take a complete and saddle-free GMN differential ϕ on a Riemann surface S which defines a point of the space $\text{Quad}(\mathbb{S}, \mathbb{M})_0$. The basic link with the combinatorics of ideal triangulations is the following.

Lemma 10.1. *Taking one generic trajectory from each horizontal strip of ϕ defines an ideal triangulation $T(\phi)$ of the surface (\mathbb{S}, \mathbb{M}) , well-defined up to the action of $\text{MCG}(\mathbb{S}, \mathbb{M})$.*

Proof. Let us identify (\mathbb{S}, \mathbb{M}) with the marked bordered surface associated to (S, ϕ) . This identification is unique up to the action of the group of orientation-preserving diffeomorphisms of (\mathbb{S}, \mathbb{M}) . In each horizontal strip h for ϕ choose a corresponding generic trajectory g_h . Note that if g_h tends to a pole p of order $m+2$, then it approaches p along one of the m distinguished tangent vectors at p . It therefore defines a path δ_h in the surface \mathbb{S} connecting two (not necessarily distinct) points of \mathbb{M} , which we denote δ_h .

The different δ_h are clearly non-intersecting in their interiors, and by Lemma 3.2 there are the correct number n of them. The fact that they are arcs corresponds to the statement that the original separating trajectories g_h are not contractible relative to their endpoints through paths with interiors in $S \setminus \text{Crit}_\infty(\phi)$. This follows from the fact that they are minimal geodesics. \square

The triangulation $T(\phi)$ is called the WKB triangulation in [14]. By definition there is a bijection $e \mapsto h_e$ between the edges of $T(\phi)$ and the horizontal strips of ϕ .

Lemma 10.2. *Under the bijection $e \mapsto h_e$ an edge e of $T(\phi)$ is self-folded precisely if the corresponding horizontal strip h_e is degenerate.*

Proof. This is clear from Figure 12: the two zeroes in the boundary of a non-degenerate strip h_e are distinct, so there are four neighbouring strips, whose corresponding edges form the two triangles containing e ; in the case of a degenerate strip h_f there is a unique neighbouring strip, necessarily non-degenerate, corresponding to the unique encircling edge e of the self-folded triangle containing the self-folded edge f . \square

Note that for any puncture $p \in \mathbb{P}$ the residue $\text{Res}_p(\phi)$ is not real, since a double pole with a real residue is contained in a degenerate ring domain, whose boundary consists of saddle connections. Suppose that we fix a signing of ϕ as in Section 6.2; this consists of a choice of sign for the residue $\text{Res}_p(\phi)$ at each puncture $p \in \mathbb{P}$. The WKB triangulation $T = T(\phi)$ then also has a naturally defined signing $\epsilon = \epsilon(\phi)$: for a puncture $p \in \mathbb{P}$ we define $\epsilon(p) \in \{\pm 1\}$ by the condition

$$(10.1) \quad \epsilon(p) \cdot \text{Res}_p(\phi) \in \mathfrak{h},$$

where $\mathfrak{h} \subset \mathbb{C}$ is the upper half-plane. We refer to $(T(\phi), \epsilon(\phi))$ as the *signed WKB triangulation* of the signed differential ϕ .

10.2. Hat-homology and the edge lattice. Let (\mathbb{S}, \mathbb{M}) be a marked bordered surface, and take a complete and saddle-free differential $\phi \in \text{Quad}(\mathbb{S}, \mathbb{M})$. If e is an edge of the WKB triangulation $T = T(\phi)$ we denote by $\alpha_e = \alpha_{h_e}$ the standard saddle class of the corresponding horizontal strip h_e . The edge lattice $\Gamma(T)$ introduced in Section 8.4 then has the following geometric interpretation.

Lemma 10.3. *There is an isomorphism of abelian groups*

$$\mu: \Gamma(T) \rightarrow \hat{H}(\phi),$$

such that for each edge $e \in T$, the basis element $[e]$ is mapped to the standard saddle class α_e of the corresponding horizontal strip h_e . This map takes the form $\langle -, - \rangle$ to the intersection form.

Proof. The map μ is an isomorphism by Lemma 3.2, since it takes a basis to a basis. We must just check the relation

$$\alpha_e \cdot \alpha_f = c(f, e) - c(e, f)$$

for all edges e and f of the triangulation T . Examining Figure 11 it is clear that α_e meets α_f precisely if e and f are two sides of the same face of T . The intersection then occurs at the unique point of the spectral cover lying over the zero of the differential contained in this face.




FIGURE 31. The oriented foliation on the spectral cover above a simple zero.

In the case when e and f are not self-folded, a glance at Figure 31 shows that the intersection of two such cycles is ± 1 depending on whether e and f occur in clockwise or anticlockwise order. If f is self-folded and e is the encircling edge of the corresponding self-folded triangle, then $\alpha_e \cdot \alpha_f = 0$, since α_e meets α_f twice with opposite signs. Comparing with (8.2) we see that the claimed relation holds also in this case. \square

In Section 8.4 we also considered a modified basis $\{e\}$ of the edge lattice $\Gamma(T)$, indexed by the edges of T . Let us define classes

$$\gamma_e = \mu(\{e\}) \in \hat{H}(\phi),$$

where μ is the map of Lemma 10.3. We can now give some geometric justification for these classes γ_e , and hence also the basis $\{e\}$.

As explained in the proof of Lemma 10.2, any degenerate horizontal strip h_f is enclosed by a non-degenerate strip h_e , and taking a generic trajectory from each strip then gives a self-folded triangle, with self-folded edge f and enclosing edge e . Note that the standard saddle connection in the strip h_f is closed and lifts to a singular curve in \hat{S} which is a bouquet of two circles. By definition of the basis $\{e\} \in \Gamma(T)$, we have

$$(\gamma_e, \gamma_f) = (\alpha_e, \alpha_e + \alpha_f).$$

These classes are both represented by simple closed curves in \hat{S} obtained by lifting the paths illustrated by dotted arcs in the two sides of Figure 33.

10.3. Flips and pops. We can lift the stratification of Section 5.2 to the étale cover of complete, signed differentials

$$\text{Quad}^\pm(\mathbb{S}, \mathbb{M})_0 \rightarrow \text{Quad}(\mathbb{S}, \mathbb{M})_0.$$

We call the connected components of B_0 chambers. The signed WKB triangulation is constant in each chamber.

The points of the locally-closed subset $F_2 = B_2 \setminus B_0$ consist of complete, signed differentials with a single saddle trajectory γ . We think of the connected components of F_2 as walls. Let us now consider the behaviour of the signed WKB triangulation as we cross such a wall.

Suppose then that $\phi_0 \in F_2 \subset \text{Quad}^\pm(\mathbb{S}, \mathbb{M})_0$ is a complete signed GMN differential lying on a wall, with a unique saddle trajectory γ . By Proposition 5.5 we can find $r > 0$ such that for all $0 < t \leq r$ the signed differentials

$$\phi_\pm(t) = e^{\pm it} \cdot \phi_0$$

are saddle-free and complete. Consider the WKB triangulations $T_\pm = T(\phi_\pm(r))$ with their signings ϵ_\pm . Note that the wall has a natural orientation to it: we make the convention that as we cross from ϕ_- to ϕ_+ the period $Z_\phi(\hat{\gamma})$ moves in a clockwise direction around the origin.

There is an isomorphism $F: \Gamma(T_-) \rightarrow \Gamma(T_+)$ defined by the following commutative diagram

$$(10.2) \quad \begin{array}{ccc} \Gamma(T_-) & \xrightarrow{F} & \Gamma(T_+) \\ \mu_- \downarrow & & \downarrow \mu_+ \\ \hat{H}(\phi_-) & \xrightarrow{\text{GM}} & \hat{H}(\phi_+) \end{array}$$

where the bottom arrow is given by the Gauss-Manin connection.

Proposition 10.4. *Take a complete signed GMN differential $\phi_0 \in F_2$ with a unique saddle trajectory γ . There are two possible cases.*

- (a) *The ends of the saddle trajectory γ are distinct. Then the signed triangulations (T_\pm, ϵ_\pm) are related by a flip in a non-self-folded edge e . Identifying the edges of these triangulations in the standard way, the map F is given by*

$$F([f]) = [f] + c(e, f) [e],$$

for all edges f , where $c(-, -)$ is computed in the triangulation T_+ , and we set $c(e, e) = -2$.

- (b) *The saddle trajectory γ is closed and forms the boundary of a degenerate ring domain centered on a double pole p with real residue. The triangulations $T_- = T_+$ are equal, and the pole p is the centre of a self-folded triangle with encircling and self-folded edges e and f respectively. The signed triangulations (T_\pm, ϵ_\pm) are related*

by a pop at p , and the map F exchanges the basis elements $\{e\}$ and $\{f\}$, leaving all other elements of the basis fixed.

Proof. Case (a) is illustrated in Figure 32 (the four poles represented by the black dots need not be distinct on the surface, however). The central picture represents ϕ_0 with




FIGURE 32. The separating trajectories on either side of a flip wall.

a single saddle trajectory appearing in the boundary of two neighbouring horizontal strips or half-planes. The wall-crossing is effected by rotating ϕ_0 . Thus to find out what happens to the saddle trajectory, we just need to consider trajectories of small nonzero phase in the two horizontal strips or half-planes. The result is as illustrated on the two sides of the figure. The associated triangulations T_{\pm} are related by a flip in a non-self-folded edge e , exactly as shown in Figure 30.

Identifying the edges of the two triangulations T_{\pm} via the obvious bijection, we see from the picture that the Gauss-Manin connection satisfies

$$\text{GM}(\alpha_f) = \alpha_f + c(e, f) \alpha_e$$

for all edges f , where $c(-, -)$ is computed in the triangulation T_+ , and we set $c(e, e) = -2$.




FIGURE 33. The separating trajectories on either side of a pop wall. The hat-homology classes γ_e are represented by the lifts of the dotted arcs.

For (b) note that if γ is closed then it is necessarily the boundary of a ring domain, which must be degenerate because there is only one saddle trajectory. This case is illustrated in Figure 33. The central picture again represents ϕ_0 with its degenerate

ring domain encased in a horizontal strip. The wall-crossing is effected by rotating the differential, so to find the separating trajectories on either side of the wall it is enough to consider trajectories of small nonzero phase. The result on either side of the wall is a degenerate horizontal strip h_f , encased in a non-degenerate horizontal strip h_e .

The WKB triangulations T_{\pm} are the same, with a self-folded triangle with self-folded edge f and encircling edge e . The hat-homology class $\hat{\gamma} = \alpha_f$ is equal to the residue class β_p , where p is the double pole at the centre of the degenerate ring domain. Since $Z_{\phi}(\hat{\gamma})$ crosses the real axis as ϕ crosses the wall, the signing $\epsilon(p)$ given by (10.1) changes, and the two signed WKB triangulations on either side of the wall are related by a pop. The two classes $\{\gamma_e, \gamma_f\} = \{\alpha_e, \alpha_e + \alpha_f\}$ are the same on both sides of the wall, but their labelling by the edges e, f is exchanged (see the dotted arcs in Figure 33). It follows that F exchanges the two basis elements $\{e\}$ and $\{f\}$ as claimed. \square

10.4. Stability conditions from saddle-free differentials. Let us now assume that our marked bordered surface (S, \mathbb{M}) is amenable, and take a saddle-free, complete, signed differential

$$\phi \in B_0 \subset \text{Quad}^{\pm}(S, \mathbb{M})_0.$$

Let (T, ϵ) denote the signed WKB triangulation of ϕ , and consider the category $\mathcal{D}(T, \epsilon)$ with its canonical heart $\mathcal{A}(T, \epsilon)$. The simple objects $S_e \in \mathcal{A}(T, \epsilon)$ are naturally indexed by the edges of the triangulation T , and so too are the classes γ_e .

Lemma 10.5. *There is an isomorphism of abelian groups*

$$\nu: K(\mathcal{D}(T, \epsilon)) \rightarrow \hat{H}(\phi),$$

taking the class of a simple object S_e to the corresponding class γ_e , and taking the Euler form to the intersection form.

Proof. This is immediate by combining Lemma 9.10 and Lemma 10.3. \square

The following result gives the basic link between quadratic differentials and stability conditions.

Lemma 10.6. *There is a unique stability condition $\sigma(\phi) \in \text{Stab } \mathcal{D}(T, \epsilon)$ whose heart is the standard heart $\mathcal{A}(T, \epsilon) \subset \mathcal{D}(T, \epsilon)$, and whose central charge*

$$Z: K(\mathcal{D}(T, \epsilon)) \rightarrow \mathbb{C}$$

corresponds to the period of ϕ under the isomorphism of Lemma 10.5.

Proof. Define Z via the isomorphism ν of Lemma 10.5. Then for each edge e the corresponding central charge $Z(S_e) = Z_{\phi}(\gamma_e)$ lies in the upper half-plane. Indeed, by the definition of the basis $\{e\}$ in Section 8.4, the classes $\gamma_e = \mu(\{e\})$ are positive linear combinations of the classes $\alpha_e = \mu([e])$, whose periods lie in the upper half-plane by

definition. Since the standard heart $\mathcal{A}(T, \epsilon)$ is of finite length, this is enough to give a stability condition. \square

10.5. Wall-crossing. We now describe how the stability conditions $\sigma(\phi)$ defined in the last section behave as the differential ϕ crosses walls in $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})_0$ of the sort considered in Section 10.3. Consider a complete, signed GMN differential

$$\phi_0 \in F_2 \subset \text{Quad}^\pm(\mathbb{S}, \mathbb{M})_0$$

with a unique saddle trajectory γ . Take $r > 0$ such that for all $0 < t \leq r$ the signed differentials

$$\phi_\pm(t) = e^{\pm it} \cdot \phi_0$$

are saddle-free and complete. Consider the corresponding signed WKB triangulations $(T_\pm, \epsilon_\pm) = (T(\phi_\pm(r), \epsilon(\phi_\pm(r)))$ and their associated categories

$$\mathcal{D}_\pm = \mathcal{D}(T_\pm, \epsilon_\pm)$$

with their standard hearts $\mathcal{A}_\pm = \mathcal{A}(T_\pm, \epsilon_\pm)$. For $0 < t < r$ we set

$$\sigma_\pm(t) = \sigma(\phi_\pm(t)) \in \text{Stab}(\mathcal{D}_\pm).$$

The following result shows that these different stability conditions glue together in the appropriate way.

Proposition 10.7. *There is a canonical equivalence $\Psi: \mathcal{D}_- \cong \mathcal{D}_+$ with the following two properties*

(a) *the diagram of isomorphisms*

$$\begin{array}{ccc} K(\mathcal{D}_-) & \xrightarrow{\Psi} & K(\mathcal{D}_+) \\ \nu_- \downarrow & & \downarrow \nu_+ \\ \hat{H}(\phi_-) & \xrightarrow{\text{GM}} & \hat{H}(\phi_+) \end{array}$$

commutes, where the bottom arrow is given by the Gauss-Manin connection, and the vertical arrows are the isomorphisms of Lemma 10.5;

(b) *the stability conditions $\Psi(\sigma_-(t))$ and $\sigma_+(t)$ on \mathcal{D}_+ become arbitrarily close as $t \rightarrow 0$.*

Proof. According to Proposition 10.4 there are two cases, the flip and the pop.

In the first case, the signed triangulations (T_\pm, ϵ_\pm) are related by a flip in an edge e , and we take Ψ to be the equivalence Φ_+ of Theorem 9.5. Part (a) then follows by comparing the formulae of Lemma 9.11 and Proposition 10.4 using Lemma 9.13. To prove (b), note that Theorem 9.5 shows that the heart of the stability condition $\Psi(\sigma_-(t))$ is the tilted heart $\mu_{S_e}^+(\mathcal{A}_+)$. By Lemma 7.9 the regions in $\text{Stab}(\mathcal{D}_+)$ consisting of stability conditions with hearts \mathcal{A}_+ and $\mu_{S_e}^+(\mathcal{A}_+)$ are glued together along a common boundary component to make a larger region on which the period map is still injective. Since part (a) shows

that the central charges of the two given stability conditions approach one another, the result follows.

In the case of the pop, the signed triangulations $(T_{\pm}, \epsilon_{\pm})$ differ by a pop, and we take Ψ to be equivalence of Theorem 9.6. Part (a) then follows by comparing the formulae of Lemma 9.12 and Proposition 10.4. To prove (b) note that since Ψ preserves the canonical hearts, all the stability conditions $\sigma_{\pm}(t)$ have the same heart, and since their central charges approach each other, they become arbitrarily close. \square

11. PROOFS OF THE MAIN RESULTS

In this section we prove our main results. Throughout (S, \mathbb{M}) is a fixed marked bordered surface. For the first five sections we shall assume that (S, \mathbb{M}) is amenable. In Section 11.6 we shall examine what can be said without this assumption.

11.1. General set-up. Let us fix a free abelian group Γ of rank n . We consider the space of framed differentials $\text{Quad}^{\Gamma}(S, \mathbb{M})$. A point of this space corresponds to a GMN differential ϕ on a Riemann surface S , equipped with a framing of the extended hat-homology group

$$\theta: \Gamma \cong \hat{H}^e(\phi)$$

as in Section 6.5. By abuse of notation, we will simply write $\phi \in \text{Quad}^{\Gamma}(S, \mathbb{M})$.

Let us fix a base-point

$$\phi_0 \in \text{Quad}^{\Gamma}(S, \mathbb{M}),$$

which we may as well assume is complete, saddle-free and generic. Recall that the space of framed differentials on (S, \mathbb{M}) is not usually connected, so we define

$$\text{Quad}_{*}^{\Gamma}(S, \mathbb{M}) \subset \text{Quad}^{\Gamma}(S, \mathbb{M})$$

to be the connected component containing ϕ_0 .

Let us choose a signing for the differential ϕ_0 , as in Section 6.2. We claim that any point in $\text{Quad}_{*}^{\Gamma}(S, \mathbb{M})$ then also has a canonical signing. To see this, note that to specify a signing of a differential ϕ is to specify a choice of sign for the residue class $\beta_p \in \hat{H}(\phi)$ at each double pole p of ϕ . Given a framing of such a differential, the classes β_p correspond to fixed classes in Γ . It follows that if we choose a sign for the β_p at the base-point ϕ_0 , then this sign propagates throughout $\text{Quad}_{*}^{\Gamma}(S, \mathbb{M})$.

We want to study stability conditions on the CY_3 triangulated category $\mathcal{D}(S, \mathbb{M})$. More precisely, let

$$(T_0, \epsilon_0) = (T(\phi_0), \epsilon(\phi_0))$$

be the signed WKB triangulation associated to the signed differential ϕ_0 , and define $\mathcal{D} = \mathcal{D}(T_0, \epsilon_0)$. We identify the Grothendieck group $K(\mathcal{D})$ with Γ using the isomorphism

$$(11.1) \quad \nu_0^{-1} \circ \theta_0: \Gamma \rightarrow K(\mathcal{D}),$$

obtained by composing the framing θ_0 with the inverse of the map of Lemma 10.5.

The distinguished connected component $\text{Tilt}_\Delta(\mathcal{D}) \subset \text{Tilt}(\mathcal{D})$ of the tilting graph of \mathcal{D} is the one containing the standard heart $\mathcal{A}(T_0, \epsilon_0)$. As explained in Section 7.7, there is a corresponding distinguished connected component

$$\text{Stab}_\Delta(\mathcal{D}) \subset \text{Stab}(\mathcal{D})$$

of the space of stability conditions on \mathcal{D} . The group of reachable autoequivalences $\text{Aut}_\Delta(\mathcal{D}) \subset \text{Aut}(\mathcal{D})$ preserves this connected component. We define

$$\text{Aut}_\Delta^0(\mathcal{D}) \subset \text{Aut}_\Delta(\mathcal{D})$$

to be the subgroup of reachable autoequivalences which act by the identity on $K(\mathcal{D})$.

Remark 11.1. *Later, as a consequence of Corollary 11.12, we will see that*

- (a) *the group $\text{Aut}_\Delta(\mathcal{D})$ is precisely the group of autoequivalences preserving the connected component $\text{Stab}_\Delta(\mathcal{D})$,*
- (b) *the group $\text{Aut}_\Delta^0(\mathcal{D})$ is non-trivial, and in fact contains all even powers of the shift functor.*

The subgroup $\text{Nil}_\Delta(\mathcal{D}) \subset \text{Aut}_\Delta(\mathcal{D})$ of negligible autoequivalences is precisely the subgroup of elements acting trivially on $\text{Stab}_\Delta(\mathcal{D})$. The quotient group

$$\mathcal{A}ut_\Delta(\mathcal{D}) = \text{Aut}_\Delta(\mathcal{D}) / \text{Nil}_\Delta(\mathcal{D})$$

therefore acts effectively. Negligible autoequivalences fix the simple objects of hearts $\mathcal{A} \in \text{Tilt}_\Delta(\mathcal{D})$, and hence act trivially on $K(\mathcal{D})$, so we can also form the quotient

$$\mathcal{A}ut_\Delta^0(\mathcal{D}) = \text{Aut}_\Delta^0(\mathcal{D}) / \text{Nil}_\Delta(\mathcal{D}).$$

In the next five sections we shall prove the following result.

Theorem 11.2. *There is an isomorphism of complex orbifolds*

$$\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) \cong \text{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta(\mathcal{D}).$$

This result implies Theorems 1.2 and 1.3 from the Introduction, except that it doesn't cover the non-amenable cases (a)–(d) listed after the statement of Theorem 1.3. These exceptional cases will be discussed in Section 11.6 below, and some of them are worked out explicitly in Section 12.

11.2. Construction of the map. In this section we construct a map from framed quadratic differentials to stability conditions.

Proposition 11.3. *There is a holomorphic map of complex manifolds K fitting into a commutative diagram*

$$(11.2) \quad \begin{array}{ccc} \mathrm{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M}) & \xrightarrow{K} & \mathrm{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta^0(\mathcal{D}) \\ & \searrow \pi & \swarrow \pi \\ & \mathrm{Hom}_\mathbb{Z}(\Gamma, \mathbb{C}) & \end{array}$$

and which commutes with the \mathbb{C} -actions on both sides.

Proof. The space on the left of the diagram (11.2) is a manifold by Proposition 6.6 (note that the surfaces listed for which this result fails are not amenable). The space on the right is a manifold because the action of the group $\mathcal{A}ut_\Delta^0(\mathcal{D})$ on the connected component $\mathrm{Stab}_\Delta(\mathcal{D})$ is free. To see this, note that if an autoequivalence Φ fixes a stability condition $\sigma \in \mathrm{Stab}_\Delta(\mathcal{D})$ and acts trivially on $K(\mathcal{D})$, then, because the period map π is a local isomorphism, it must act trivially on a neighbourhood of σ , and hence on the whole connected component $\mathrm{Stab}_\Delta(\mathcal{D})$. But then it is negligible and hence defines the identity in $\mathcal{A}ut_\Delta^0(\mathcal{D})$.

The action of \mathbb{C} on the right of (11.2) is the standard one of Section 7.5; the element $t \in \mathbb{C}$ acts at the level of central charges by $Z(E) \mapsto e^{i\pi t} \cdot Z(E)$. After the event it will follow from Corollary 11.12 that $2 \in \mathbb{C}$ acts trivially, so that this factors via a \mathbb{C}^* action, but we don't know this yet. The action of \mathbb{C} on the left is the pullback of the standard \mathbb{C}^* action rescaling the quadratic differential, via the map $\mathbb{C} \rightarrow \mathbb{C}^*$ defined by $t \mapsto e^{2\pi i t}$. This action lifts to framed differentials by continuity; note that $1 \in \mathbb{C}$ acts trivially on the underlying quadratic differential, but multiplies the framing isomorphism by -1 .

The maps π in the diagram (11.2) are both local isomorphisms. The map on the left is the standard period map on framed differentials. The map on the right sends a stability condition to its central charge, which we consider as a group homomorphism $Z: \Gamma \rightarrow \mathbb{C}$ by composing it with the isomorphism (11.1). Let $t \in \mathbb{C}$ act on the space of central charges $\mathrm{Hom}_\mathbb{Z}(\Gamma, \mathbb{C})$ via the map $Z(E) \mapsto e^{i\pi t} \cdot Z(E)$ as above. Then because periods of quadratic differentials are given by integrals of $\sqrt{\phi}$, both the maps π are \mathbb{C} -equivariant.

As soon as we know that K is continuous it is automatically holomorphic, and in fact, a local isomorphism. The \mathbb{C} -equivariance is also automatic, just because each of the two maps π is a local isomorphism. We first define a map

$$K_0: B_0 \rightarrow \mathrm{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta^0(\mathcal{D}).$$

The key point is to then use the stratification

$$B_0 \subset B_1 \subset B_2 \subset \cdots \subset B_k \subset \cdots \subset \mathrm{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M})_0 \subset \mathrm{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M})$$

and inductively extend the domain of definition of the map across larger strata.

Let ϕ be a saddle-free complete framed differential defining a point in $\text{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M})$. Recall that the signed WKB triangulation $(T(\phi), \epsilon(\phi))$ is well-defined up to the action of the mapping class group $\text{MCG}(\mathbb{S}, \mathbb{M})$. Take an equivalence

$$\Psi: \mathcal{D}(T(\phi), \epsilon(\phi)) \longrightarrow \mathcal{D}(T(\phi_0), \epsilon(\phi_0))$$

with the following two properties

- (i) Ψ maps the distinguished connected components of the tilting graphs of $\mathcal{D}(T(\phi), \epsilon(\phi))$ and $\mathcal{D}(T(\phi_0), \epsilon(\phi_0))$ one to the other;
- (ii) the map ψ on Grothendieck groups induced by Ψ makes the following diagram commute

$$(11.3) \quad \begin{array}{ccc} K(\mathcal{D}(T(\phi), \epsilon(\phi))) & \xrightarrow{\psi} & K(\mathcal{D}(T(\phi_0), \epsilon(\phi_0))) \\ \nu \downarrow & & \downarrow \nu_0 \\ \hat{H}(\phi) & & \hat{H}(\phi_0) \\ & \swarrow \theta \quad \searrow \theta_0 & \\ & \Gamma & \end{array}$$

where ν and ν_0 are the isomorphisms of Lemma 10.5, and θ and θ_0 are the framing isomorphisms.

To see that such an equivalence Ψ exists, connect ϕ_0 to ϕ by some path in $\text{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M})$. Since the subset of differentials with simple poles is locally cut out by complex hyperplanes we can assume that this path lies in $\text{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M})_0$. By Corollary 5.8 and the fact that $\text{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M})_0$ is a covering space of $\text{Quad}(\mathbb{S}, \mathbb{M})_0$, we can then deform the path so that it lies in B_2 and has only finitely many points in F_2 . Applying Proposition 10.7 to each of these points, and taking the composite of the given equivalences gives a suitable equivalence Ψ .

Any two such equivalences Ψ differ by a reachable autoequivalence acting trivially on $K(\mathcal{D})$, so we obtain a well-defined map K_0 by setting

$$K_0(\phi) = \Psi(\sigma(\phi)) \in \text{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta^0(\mathcal{D}).$$

The diagram (11.2) then commutes by definition.

Our task is now to successively lift K to the inverse images of the various strata B_i . Let us assume inductively that K is defined and continuous on the open subset B_{p-1} . As remarked above, K is invariant under small rotations, i.e.

$$K(e^{i\pi\theta} \cdot \phi) = e^{i\pi\theta} \cdot K(\phi) \text{ for } 0 < |\theta| \ll 1,$$

because both maps π of (11.2) are local isomorphisms. Take a point $\phi_0 \in F_p$. By Proposition 5.5 there is some $r > 0$ such that

$$(11.4) \quad 0 < |t| < r \implies e^{it} \cdot \phi_0 \in B_{p-1}.$$

By the \mathbb{C} -equivariance property, the limits

$$\sigma_{\pm}(\phi_0) = \lim_{t \rightarrow 0^+} K(e^{\pm it} \cdot \phi_0)$$

both exist, and the diagram (11.2) shows that they have the same central charge. The point is to show that they are equal.

When $p = 2$ this involves extending from saddle-free differentials to differentials lying on a single wall, and the result follows from Proposition 10.7. Thus we can assume that $p > 2$. Note that the stability conditions $\sigma_{\pm}(\phi_0)$ vary continuously on F_p , by the fact that their central charges do, and using the remark following Proposition 5.5. Thus the question of whether they are equal has a constant answer on each connected component of F_p . The result then follows from Proposition 5.7.

The final step is to extend K across the incomplete locus. Suppose that $\phi \in \text{Quad}_{*}^{\Gamma}(\mathbb{S}, \mathbb{M})$ has simple poles, and fix $0 < \epsilon < \frac{1}{8}$. Using Proposition 6.7 we can find complete, generic differentials ψ arbitrarily close to ϕ , and such that

$$|Z_{\phi}(\gamma) - Z_{\psi}(\gamma)| < \epsilon |Z_{\psi}(\gamma)|,$$

for all classes $\gamma \in \Gamma$ represented by a non-closed saddle connection in ψ . Lemma 11.4 below then shows that this inequality holds for all classes represented by stable objects in the stability condition $K(\psi)$. The deformation result Proposition 7.5 then shows that K extends uniquely over ϕ . \square

11.3. Saddle trajectories and stable objects. In this section and the next we relate saddle trajectories for a generic GMN differential to the existence of stable objects in the corresponding stability condition. Let

$$\phi \in \text{Quad}_{*}^{\Gamma}(\mathbb{S}, \mathbb{M})_0$$

is a complete, framed differential, and let $\sigma = K(\phi)$ be the corresponding stability condition on \mathcal{D} , well-defined up to the action of the group $\text{Aut}_{\Delta}^0(\mathcal{D})$. Note that any saddle connection γ for ϕ has a well-defined hat-homology class in $\hat{H}(\phi)$, which we can view as an element of Γ using the framing isomorphism. Similarly, every object $E \in \mathcal{D}$ has a well-defined class in Γ using the identification (11.1).

We shall start with the following simple result which was used in the final step of the proof of Proposition 11.3 above.

Lemma 11.4. *If σ has a stable object E of class $\alpha \in \Gamma$ then ϕ has a saddle trajectory of class proportional to α .*

Proof. Rotating we can assume that E is of phase 1. Consider a generic differential ψ close to ϕ such that $Z_\psi(\alpha)$ remains real. Proposition 7.6(a) shows that the object E remains stable in the corresponding stability condition $K(\psi)$. By construction therefore ψ cannot be saddle-free, and has at least one saddle trajectory C , which by genericity must have class proportional to α . Since the class of a saddle connection has divisibility at most 2, the length of this saddle is at most $|Z_\psi(2\alpha)|$.

Applying Theorem 4.1, we can now find a sequence of such differentials ψ_i , converging to ϕ , such that the corresponding curves C_i limit to some curve C . Then C is a geodesic in ϕ which must be a union of saddle trajectories. \square

For each class $\alpha \in \Gamma$ we can consider the moduli space $\mathcal{M}_\sigma(\alpha)$ of σ -stable objects in \mathcal{D} which have class $\alpha \in \Gamma$ and phase in the interval $(0, 1]$. It is necessary to constrain the phase, since otherwise all shifts of a given stable object would have to be parameterized.

Lemma 11.5. *The moduli space $\mathcal{M}_\sigma(\alpha)$ is represented by a quasi-projective scheme.*

Proof. By rotation, we can assume that ϕ is saddle-free and therefore defines some signed triangulation (T, ϵ) . The heart of the stability condition σ is then equivalent to the category of finite-dimensional modules for the complete Jacobi algebra of the corresponding quiver with potential. This algebra $J(T, \epsilon)$ is known to be finite-dimensional¹⁰ [26, Corollary 12.6]. The moduli space $\mathcal{M}_\sigma(\alpha)$ can therefore be identified with the moduli space of θ -stable representations of $J(T, \epsilon)$ with fixed dimension vector. The claim then follows from the results of King [21]. \square

Recall the notion of a 0-generic differential from Section 5.2. Note that any such differential is, in particular, complete. In this and the next section we shall prove the following precise correspondence, which implies Theorem 1.4 from the Introduction.

Theorem 11.6. *Assume that ϕ is 0-generic and take a class $\alpha \in \Gamma$ satisfying $Z_\phi(\alpha) \in \mathbf{R}$. Then each connected component of the moduli scheme $\mathcal{M}_\sigma(\alpha)$ is either a point or a copy of \mathbb{P}^1 . Moreover*

- (a) *the zero-dimensional components of $\mathcal{M}_\sigma(\alpha)$ are in bijection with the non-closed saddle trajectories for ϕ of class α ;*
- (b) *the one-dimensional components of $\mathcal{M}_\sigma(\alpha)$ are in bijection with the non-degenerate ring domains for ϕ of class α .*

We first prove the result under an additional assumption; the general case will be dealt with in the next section.

¹⁰This could also be deduced from Corollary 11.12 below together with an argument of Nagao [19, Theorem 5.4].

Proposition 11.7. *Take assumptions as in Theorem 11.6. Suppose moreover that ϕ has at most one saddle trajectory. Then the conclusion of Theorem 11.6 holds.*

Proof. Suppose that ϕ has a unique saddle trajectory γ of some class $\alpha \in \Gamma$. Since the surface (\mathbb{S}, \mathbb{M}) is assumed to be amenable, it is not a closed surface with a single puncture, and it follows that there can be no spiral domains. Thus $\phi \in F_2$ and we are in the situation of Proposition 10.4.

Assume first that γ has distinct ends, as in Proposition 10.4(a). Small rotations of ϕ are saddle-free, and γ is represented by a standard saddle class. The corresponding stability conditions have a unique simple object S of class α . This is, in particular, stable and spherical, and all semistable objects of class proportional to α are of the form $S^{\oplus k}$, and hence strictly semistable for all $k > 1$.

Suppose instead that γ is a closed saddle trajectory. This is the situation of Proposition 10.4(b). By Proposition 10.7, stability conditions on either side of the wall obtained by small rotations of σ have the same heart. This implies that $\mathcal{P}(0) = (0)$ and so there are no semistable objects with class a multiple of α .

For the converse, suppose that the stability condition σ has a stable object of class α . Then, by construction of the map K , the differential ϕ cannot be saddle-free. Thus ϕ has a saddle trajectory γ of some class β proportional to α . Applying what we have proved in the first part, it follows that σ has at most one stable object of class β , and no stable objects of any other class proportional to β . We conclude that $\alpha = \beta$ and so ϕ has a (necessarily unique) saddle trajectory of class α . \square

11.4. Ring-shrinking again. In this section we complete the proof of Theorem 11.6. It will be convenient to denote the differential and corresponding stability condition of the statement by ϕ_+ and σ_+ respectively. Thus we consider a complete, framed, 0-generic differential

$$\phi_+ \in \text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})_0,$$

and let $\sigma_+ = K(\phi_+)$ be the corresponding stability condition on \mathcal{D} , well-defined up to the action of the group $\text{Aut}_\Delta^0(\mathcal{D})$. We may assume that ϕ_+ has more than one saddle trajectory, since otherwise we are in the situation of Proposition 11.7.

The surface (\mathbb{S}, \mathbb{M}) is amenable, hence not a once-punctured torus, so according to Section 5.9 there are two possible cases, labelled (J1) and (J2). Shrinking the unique ring domain A as in Section 5.9 gives a smooth path in $\text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})_0$ ending at a non-generic point

$$\phi \in \text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})_0$$

with either 2 or 3 saddle trajectories γ_i . Label the saddle trajectories γ_i exactly as in Section 5.9, and write $\alpha_i \in \Gamma$ for the corresponding hat-homology classes. We recall

that there is a linear relation

$$\alpha_3 = \alpha_1 + \alpha_2.$$

Our strategy will be to first understand the stable objects of phase 1 in the stability condition $\sigma = K(\phi)$ by applying Proposition 11.7 to nearby points on the other side of the wall

$$\operatorname{Im} Z_{\phi_-}(\alpha_1)/Z_{\phi_-}(\alpha_2) = 0.$$

We will then follow this information back through the ring-shrinking operation.




FIGURE 34. The quivers relevant to the two cases (J1) and (J2).

Let $\mathcal{A} = \mathcal{P}(1)$ denote the abelian category of semistable objects of phase 1 in the stability condition $\sigma = K(\phi)$. This category has finite-length, so we can model it by the category of representations of the Jacobi algebra of a quiver with potential (Q, W) . The relevant quivers Q in the two cases (J1) and (J2) are as shown in Figure 34; in both cases the potential is necessarily zero.

Proposition 11.8. *The category \mathcal{A} is equivalent to the category of finite-dimensional representations of the corresponding quiver Q .*

Proof. For definiteness we consider the more difficult case (J2).

Let n_{ij} denote the number of arrows in the quiver Q connecting vertex i to vertex j . To prove the Lemma we must show

- (a) the category \mathcal{A} has exactly 3 stable objects S_i ;
- (b) these objects have classes $[S_i] = \alpha_i$ respectively;
- (c) For all i, j we have $\dim_k \operatorname{Ext}_{\mathcal{A}}^1(S_i, S_j) = n_{ij}$.

It is easy to see that after contracting the ring domain in Figure 18 the intersection multiplicities $\alpha_i \cdot \alpha_j$ coincide with the expressions $n_{ji} - n_{ij}$. So for the last part it will be enough to show that each object S_i is spherical, and for each pair $i \neq j$, we have either

$$\operatorname{Ext}_{\mathcal{A}}^1(S_i, S_j) = 0 \text{ or } \operatorname{Ext}_{\mathcal{A}}^1(S_j, S_i) = 0.$$

Take a class $\beta \in \Gamma$ and consider the wall-and-chamber decomposition of $\operatorname{Stab}(\mathcal{D})$ with respect to the class β . Take a chamber containing σ in its closure, and containing points

$\sigma_- = K(\phi_-)$ which satisfy

$$\operatorname{Im} Z_{\phi_-}(\alpha_1)/Z_{\phi_-}(\alpha_2) < 0, \quad \operatorname{Im} Z_{\phi_-}(\beta) = 0.$$

Let us choose such a point σ_- in this chamber, and assume further that the corresponding differential ϕ_- is generic, and that it lies in the open subset U of Proposition 5.10. Of course, our choice of $\phi_- = \phi_-(\beta)$ will depend on the class β we started with.

The genericity condition implies that all saddle trajectories for ϕ_- have classes which are multiples of β . Proposition 5.10 then implies that any such saddle trajectory is one of the γ_i . Since the classes α_i are pairwise non-proportional, it follows in particular that ϕ_- has at most one saddle trajectory. Thus Proposition 11.7 applies to ϕ_- , and shows that ϕ_- has a saddle trajectory of class β precisely if σ_- has a stable object of class β .

In the case when $\beta = \alpha_i$ for some i , Proposition 11.7 implies that $\sigma_- = \sigma_-(\alpha_i)$ has a unique stable object S_i of class β , which is moreover spherical. By Proposition 7.8, the object S_i is then semistable in σ , and hence lies in the category $\mathcal{A} = \mathcal{P}(1)$.

Conversely, suppose that an object $E \in \mathcal{A}$ is stable in σ of phase 1. Proposition 7.6(a) shows that E is stable in all nearby stability conditions, and in particular we can assume that this is the case for $\sigma_- = \sigma_-(\beta)$. Then ϕ_- must have a saddle trajectory of class β , and by Proposition 5.10 it follows that $\beta = \alpha_i$ for some i . Since S_i was the unique stable object in $\sigma_-(\alpha_i)$ of class α_i it follows that $E = S_i$.

Thus the set of stable objects in $\mathcal{P}(1)$ is some subset of the objects $\{S_1, S_2, S_3\}$. It follows from this that the objects S_1 and S_2 must actually be stable in σ . For example, if S_2 was unstable, it would have a filtration by the objects S_1 and S_3 , which is impossible because α_2 is not a positive linear combination of the classes α_1 and $\alpha_3 = \alpha_1 + \alpha_2$.

Suppose there are extensions in two directions between S_1 and S_2 . Then on both sides of the wall $\operatorname{Im} Z_{\phi}(\alpha_1)/Z_{\phi}(\alpha_2) = 0$ there would be stable objects of class $[S_1] + [S_2]$, which is not the case for σ_- . It follows that the object S_3 must also be stable, since the only other possibility is that there is a sequence

$$0 \longrightarrow S_2 \longrightarrow S_3 \longrightarrow S_1 \longrightarrow 0,$$

which is impossible since S_3 is stable on the σ_- side of the wall. Using the same argument as before, we can now check that nonzero extensions between one of the objects S_1 or S_2 and S_3 only go in one direction. This completes the proof. \square

Consider stability conditions $W: K(\mathcal{A}) \rightarrow \mathbb{C}$ on the abelian category \mathcal{A} satisfying the conditions

$$(11.5) \quad Z(S_1) + Z(S_2) \in i\mathbf{R}, \quad \operatorname{Im} Z(S_1)/Z(S_2) > 0.$$

In the case (J2), assume also that $Z(S_3) \in i\mathbf{R}$. Then

Lemma 11.9. *The set of stable objects satisfying $Z(E) \in i\mathbf{R}$ is independent of the particular choice of stability condition satisfying (11.5), and is as follows:*

- (J1) a single \mathbb{P}^1 family of objects of dimension vector $(1, 1)$;
 (J2) a single \mathbb{P}^1 family of objects of dimension vector $(1, 1, 1)$, and unique objects of dimension vectors $(1, 1, 0)$ and $(0, 0, 1)$.

Proof. In the case (a) we are considering representations of the Kronecker quiver and the result is well-known. In case (b) we are considering the affine A_2 quiver, and since indecomposable representations are completely understood in terms of the real and imaginary roots, the result is again easy. \square

Note the precise correspondence with the finite-length trajectories of ϕ_+ listed in Section 5.9. The following result then completes the proof of Theorem 11.6.

Proposition 11.10. *An object $E \in \mathcal{D}$ is stable of phase 1 in σ_+ precisely if it lies in the abelian subcategory \mathcal{A} and is stable with respect to stability conditions W as above.*

Proof. As the ring domain A shrinks we move along a path of differentials $\phi_+(t)$ for $0 \leq t \leq 1$ with $\phi_+(0) = \phi_+$ and $\phi_+(1) = \phi$. Let $\sigma_+(t) = K(\phi_+(t))$ be the corresponding stability conditions. The first claim is that the class of stable objects of phase 1 in the stability condition $\sigma_+(t)$ is constant for $0 \leq t < 1$. To prove this it will be enough to show that if E is a stable object of phase 1 in some differential $\sigma_+(t)$, then the class β of E is a multiple of the class α of the ring domain, and hence remains of phase 1 for all t . This follows immediately from Lemma 11.4 and the list of saddle trajectories for σ_+ given in Section 5.9.

Suppose that $E \in \mathcal{D}$ is stable in σ_+ of phase 1. Then by the above, the class of E is proportional to α , and moreover E is stable in $\sigma_+(t)$ for $0 \leq t < 1$. It follows that E is at least semistable in σ , and hence that $E \in \mathcal{A} = \mathcal{P}(1)$. Thus we are reduced to understanding which objects $E \in \mathcal{A}$ whose classes are proportional to α are stable in the stability conditions $\sigma_+(t)$ for $0 \leq t < 1$.

Consider the central charge of $\sigma_+(t)$ and rotate it by an angle of $\pi/2$. It then induces a stability function W on \mathcal{A} satisfying the conditions (11.5). It follows that an object $E \in \mathcal{A}$ of class proportional to α is stable in $\sigma_+(t)$ precisely if it is stable with respect to the stability conditions W of Lemma 11.9. \square

11.5. Completion of the proof. The following result will be enough to complete the proof of Theorem 11.2.

Proposition 11.11. *The map K of Proposition 11.3 is an isomorphism of complex manifolds.*

Proof. We begin by showing that K is injective; it then follows that it is an open embedding because it commutes with the period maps, which are local isomorphisms. Suppose we have two distinct framed differentials

$$\phi_1, \phi_2 \in \text{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M})$$

such that $K(\phi_2) = \Phi(K(\phi_1))$ for some autoequivalence $\Phi \in \text{Aut}_\Delta^0(\mathcal{D})$. Since the period maps are local isomorphisms, and K commutes with these maps, if we deform both the ϕ_i maintaining the condition that their periods are equal, we will also preserve the condition $K(\phi_2) = \Phi(K(\phi_1))$. Thus we can assume that the ϕ_i are saddle-free. Let (T_i, ϵ_i) be the associated signed WKB triangulations.

By definition, $K(\phi_i) = \Psi_i(\sigma(\phi_i))$, where the Ψ_i are equivalences

$$\Psi_i: \mathcal{D}(T_i, \epsilon_i) \rightarrow \mathcal{D}(T_0, \epsilon_0)$$

satisfying the conditions (i) and (ii) of the proof of Proposition 11.3. It follows that $\Psi_2^{-1} \circ \Phi \circ \Psi_1$ takes the canonical heart $\mathcal{A}(T_1, \epsilon_1) \subset \mathcal{D}(T_1, \epsilon_1)$ to the canonical heart $\mathcal{A}(T_2, \epsilon_2) \subset \mathcal{D}(T_2, \epsilon_2)$. In particular, the quivers $Q(T_i)$ are isomorphic. Thus by Proposition 8.5, the WKB triangulations T_i differ by an orientation-preserving diffeomorphism.

Examining Figure 5 it is easy to see that this implies that the differentials ϕ_i have the same horizontal strip decomposition in the sense of Section 4.5. The fact that the equivalences Ψ_i satisfy the condition (ii) of the proof of Proposition 11.3 implies further that the horizontal strips of the ϕ_i have the same labellings by elements of Γ . Proposition 4.9 then shows that $\phi_1 = \phi_2 \in \text{Quad}^\Gamma(\mathbb{S}, \mathbb{M})$.

The fact that the image of K is closed follows from Proposition 6.8. Indeed, suppose $\sigma_n \rightarrow \sigma \in \text{Stab}_\Delta(\mathcal{D})$ with $\sigma_n = K(\phi_n)$. We can assume that each ϕ_n is complete and generic since these points are dense. Theorem 11.6 implies that the lengths of non-degenerate saddle connections in ϕ_n correspond to the masses of stable objects in σ_n . Thus it will be sufficient to show that the masses of objects in the σ_n are uniformly bounded below. By continuity it is enough to check that the masses of objects in σ are bounded below, and this follows from the support property (7.2). \square

To prove Theorem 11.2 it remains to show that K descends to give an isomorphism of complex orbifolds fitting into the commutative diagram

$$(11.6) \quad \begin{array}{ccc} \text{Quad}_*^\Gamma(\mathbb{S}, \mathbb{M}) & \xrightarrow{K} & \text{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta^0(\mathcal{D}) \\ \downarrow & & \downarrow \\ \text{Quad}_\nabla(\mathbb{S}, \mathbb{M}) & \xrightarrow{J} & \text{Stab}_\Delta(\mathcal{D}) / \mathcal{A}ut_\Delta(\mathcal{D}) \end{array}$$

Since the map on the left is a covering map, to prove that K descends it is enough to check that all framings of the differential ϕ_0 give the same stability condition up to the action of $\mathcal{A}ut_\Delta(\mathcal{D})$. This is immediate from the definition of K_0 .

To prove that the resulting map J is injective, hence an isomorphism, we follow the same argument used above to show that K is injective. Suppose that $J(\phi_1) = J(\phi_2)$. By deforming the differentials ϕ_i as before, it is enough to deal with the case when the ϕ_i are complete and saddle-free. As above we conclude that the ϕ_i have the same horizontal strip decomposition. Since they also have the same periods, they are equal. \square

Corollary 11.12. *An autoequivalence of \mathcal{D} preserves the connected component $\text{Stab}_\Delta(\mathcal{D})$ precisely if it is reachable. In particular, the shift functor $[1]$ lies in the group $\text{Aut}_\Delta(\mathcal{D})$ of reachable autoequivalences.*

Proof. One implication is always true, by Lemma 7.9. Conversely, suppose that $\Phi \in \text{Aut}(\mathcal{D})$ preserves $\text{Stab}_\Delta(\mathcal{D})$. Take a smooth path connecting σ_0 and $\Phi(\sigma_0)$ in $\text{Stab}_\Delta(\mathcal{D})$. Applying K^{-1} this defines a smooth path of framed quadratic differentials. As in the proof of Proposition 11.3, we can use Proposition 5.7 to find a homotopic path consisting entirely of tame differentials, and crossing only finitely many walls. Applying Proposition 10.7 at each of these walls shows that Φ is reachable. \square

11.6. Non-amenable cases. Suppose now that (\mathbb{S}, \mathbb{M}) is a marked bordered surface which fails to be amenable because it violates one of the first four conditions of Definition 9.3. We shall make some brief comments about what can be said in these various cases.

Case (a). Suppose that (\mathbb{S}, \mathbb{M}) does not satisfy Assumption 8.1. We refer the reader to the comments made following Assumption 8.1. We cannot deal with the case of a 4-punctured sphere, but the other cases are all described explicitly in Section 12.

Case (b). Suppose that (\mathbb{S}, \mathbb{M}) is one of the three surfaces listed in Proposition 8.5 having distinct triangulations with the same associated quiver. The isomorphism of Theorem 9.8 still holds and induces an action of the group $\text{Aut}_\Delta(\mathcal{D})$ on the set of tagged triangulations. We say that a reachable autoequivalence is *allowable* if the induced action on the quotient set

$$\text{Tri}_{\bowtie}(\mathbb{S}, \mathbb{M}) / \text{MCG}^\pm(\mathbb{S}, \mathbb{M})$$

is trivial. If we replace the group $\text{Aut}_\Delta(\mathcal{D})$ of reachable autoequivalences by the subgroup $\text{Aut}_\Delta^{\text{allow}}(\mathcal{D})$ of allowable ones then Proposition 9.9 holds for (\mathbb{S}, \mathbb{M}) with the same proof.

In these cases Theorem 11.2 still holds, with the group $\text{Aut}_\Delta(\mathcal{D})$ replaced by $\text{Aut}_\Delta^{\text{allow}}(\mathcal{D})$. The proof is the same, one just needs to check at several places that certain autoequivalences are allowable. This is easily done, and we omit the details. The case of a once-punctured disc with two points on the boundary is described in detail in Example 12.4 below.

Case (c). Consider the two non-closed surfaces listed in Proposition 8.6, for which the action of the mapping class group on ordered ideal triangulations is not free. Proposition 9.9 also fails in these two cases. These surfaces are dealt with explicitly in Examples

12.2 and 12.5 below. Note that the orbifold $\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$ has non-trivial generic automorphism group \mathbb{Z}_2 , and Theorem 11.2 continues to hold if we rigidify the orbifold $\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$ by killing this group.

Case (d). We leave the case of a closed surface with a single puncture for future research, and restrict ourselves here to some sketchy comments about some of the special features of these surfaces.

The first new feature is that the graph $\text{Tri}_\bowtie(\mathbb{S}, \mathbb{M})$ has two connected components. This means that we have two potentially distinct categories $\mathcal{D}(\mathbb{S}, \mathbb{M})$, depending on the choice of sign $\epsilon = \epsilon(p)$. But by the result [7, Prop. 10.4] referred to in the proof of Theorem 9.9, the two potentials $W(T, \pm 1)$ are right-equivalent up to scale, so in fact there is a well-defined category $\mathcal{D} = \mathcal{D}(\mathbb{S}, \mathbb{M})$.

As stated, Theorem 9.8 is false for these examples, again because the graph $\text{Tri}_\bowtie(\mathbb{S}, \mathbb{M})$ is disconnected. A related issue is that Corollary 11.12 fails, and the shift functor $[1]$ is not reachable¹¹. It might therefore be best to replace the notion of a reachable heart by a *shift-reachable* heart: one for which $\mathcal{A}[i]$ lies in $\text{Tilt}_\Delta(\mathcal{D})$ for some $i \in \mathbb{Z}$. Of course, in the case of an amenable surface, the notions of reachable and shift-reachable coincide by Corollary 11.12.

One more feature in this case is that the complete Ginzburg algebra definitely does not coincide with the uncompleted version in general, see Remark 9.7. It may well be that it is more natural to consider the uncompleted Ginzburg algebra in this context.

12. EXAMPLES

In this section we consider some special cases of Theorem 1.3 corresponding to surfaces (\mathbb{S}, \mathbb{M}) of genus $g = 0$. This leads to descriptions of spaces of stability conditions on CY_3 categories associated to certain simple quivers familiar in representation theory. We adopt a less formal approach in this section, and some of the details are left for the reader.

12.1. Unpunctured discs: A_n type. Fix an integer $n \geq 2$ and let (\mathbb{S}, \mathbb{M}) be an unpunctured disc with $n+3$ points on its boundary. This corresponds to differentials on \mathbb{P}^1 with a single pole of order $n+5$. The space $\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})$ coincides with $\text{Quad}(\mathbb{S}, \mathbb{M})$ and parameterizes differentials of the form

$$\varphi(z) = p_{n+1}(z) dz^{\otimes 2}$$

¹¹To see this, note that because there are no self-folded triangles, the residue class β_p is always either a strictly positive or strictly negative linear combination of the classes of the simple objects of any heart, and this sign is constant under mutation.

where $p_{n+1}(z)$ is a polynomial of degree $n + 1$ having simple roots, considered modulo the automorphisms of \mathbb{P}^1 which fix infinity. Taking the sum of the roots to be 0 we can reduce to differentials of the form

$$\phi(z) = \prod_{i=1}^{n+1} (z - a_i) dz^{\otimes 2}, \quad \sum_{i=1}^{n+1} a_i = 0, \quad a_i \neq a_j,$$

modulo a residual action of \mathbb{Z}_{n+3} acting by rescaling z by an $(n + 3)$ rd root of unity. Thus

$$\text{Quad}(\mathbb{S}, \mathbb{M}) \cong \text{Conf}_0^{n+1}(\mathbb{C}) / \mathbb{Z}_{n+3},$$

where $\text{Conf}_0^{n+1}(\mathbb{C})$ denotes the configuration space of $n + 1$ distinct points in \mathbb{C} with centre of mass at the origin, and the group \mathbb{Z}_{n+3} acts by multiplication by $(n + 3)$ rd roots of unity.




FIGURE 35. Triangulation of a 5-gon.

The mapping class group of the surface is

$$\text{MCG}(\mathbb{S}, \mathbb{M}) = \mathbb{Z}_{n+3}$$

and coincides with the signed mapping class group. Ideal triangulations of (\mathbb{S}, \mathbb{M}) correspond to triangulations of an $(n + 3)$ -gon. The resulting quivers are orientations of a Dynkin diagram of A_n type and necessarily have zero potential.

Our main Theorem gives an isomorphism of complex orbifolds

$$\text{Conf}_0^{n+1}(\mathbb{C}) / \mathbb{Z}_{n+3} \cong \text{Stab}_\Delta(\mathcal{D}) / \text{Aut}_\Delta(\mathcal{D}).$$

Note that there is a short exact sequence

$$1 \longrightarrow \text{Br}(A_n) \longrightarrow \pi_1 \text{Quad}(\mathbb{S}, \mathbb{M}) \longrightarrow \mathbb{Z}_{n+3} \longrightarrow 1,$$

where the Artin braid group $\text{Br}(A_n)$ is the fundamental group of the configuration space $\text{Conf}_0^{n+1}(\mathbb{C})$. On the other hand, the group $\text{Sph}(\mathcal{D}) \subset \text{Aut}(\mathcal{D})$ is isomorphic to $\text{Br}(A_n)$ by [30, Theorem 1.3]. The sequence of Lemma 9.9 therefore becomes

$$1 \longrightarrow \text{Br}(A_n) \longrightarrow \text{Aut}_\Delta(\mathcal{D}) \longrightarrow \mathbb{Z}_{n+3} \longrightarrow 1.$$

The fact that these two sequences coincide suggests that $\text{Stab}_\Delta(\mathcal{D})$ is simply-connected.

Example 12.1. *The non-amenable case $n = 0$ corresponding to an unpunctured disc with 3 marked points on its boundary is very degenerate. The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ is a single point with automorphism group \mathbb{Z}_3 , corresponding to the unique differential $\phi(z) = z dz^{\otimes 2}$. The mapping class group is $\text{MCG}(\mathbb{S}, \mathbb{M}) = \mathbb{Z}_3$. There is a unique ideal triangulation, but it contains no edges, and the associated quiver is empty.*

Example 12.2. *Consider the non-amenable case $n = 1$ corresponding to an unpunctured disc with four marked points on its boundary. The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ consists of differentials of the form*

$$\phi(z) = (z^2 + c) dz^{\otimes 2}, \quad c \in \mathbb{C}^*,$$

modulo the action of \mathbb{Z}_4 acting on z by multiplication by i . Thus

$$\text{Quad}(\mathbb{S}, \mathbb{M}) \cong \mathbb{C}^* / \mathbb{Z}_4$$

with the generator of \mathbb{Z}_4 acting by change of sign. Note that the generic stabilizer group is \mathbb{Z}_2 .

There are two ideal triangulations of (\mathbb{S}, \mathbb{M}) , each with a single edge. These are related by the action of the mapping class group \mathbb{Z}_4 . Note that Proposition 8.6 breaks down in this case: the element of \mathbb{Z}_4 of order 2 fixes all triangulations.

The quiver corresponding to either triangulation has a single vertex and no arrows. The category \mathcal{D} is generated by a single spherical object S , and the group $\mathcal{A}ut_{\Delta}(\mathcal{D})$ is generated by the shift functor $[1]$. Proposition 9.9 also breaks down in this case, since $\text{Sph}(\mathcal{D}) \cong \mathbb{Z}$ is generated by the second shift $[2]$.

The quotient space $\text{Stab}(\mathcal{D}) / \langle [2] \rangle$ is isomorphic to \mathbb{C}^ with co-ordinate $Z(S)$. Thus*

$$\text{Stab}(\mathcal{D}) / \mathcal{A}ut_{\Delta}(\mathcal{D}) \cong \mathbb{C}^* / \mathbb{Z}_2,$$

with the generator of \mathbb{Z}_2 acting by change of sign. There is a morphism

$$\text{Quad}(\mathbb{S}, \mathbb{M}) \rightarrow \text{Stab}(\mathcal{D}) / \mathcal{A}ut_{\Delta}(\mathcal{D})$$

given by setting $Z(S) = \pi ic$, but it is not an isomorphism of orbifolds, since the generic automorphism groups are different.

12.2. Punctured discs: D_n type. Fix an integer $n \geq 2$ and let (\mathbb{S}, \mathbb{M}) be a once-punctured disc with n points on its boundary. This corresponds to differentials on \mathbb{P}^1 with polar type $(2, n+2)$.

The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ consists of differentials of the form

$$\phi(z) = \prod_{i=1}^n (z - a_i) \frac{dz^{\otimes 2}}{z^2}, \quad a_i \neq a_j,$$

modulo the action of \mathbb{C}^* rescaling z . Thus

$$\text{Quad}(\mathbb{S}, \mathbb{M}) \cong \text{Conf}^n(\mathbb{C}) / \mathbb{Z}_n,$$

where $\text{Conf}^n(\mathbb{C})$ denotes the configuration space of n distinct points in \mathbb{C} , and the group \mathbb{Z}_n acts by multiplication by n th roots of unity.

Note that the product $a = \prod_{i=1}^n a_i$ is invariant under the \mathbb{Z}_n action, and hence defines a map

$$a: \text{Quad}(\mathbb{S}, \mathbb{M}) \rightarrow \mathbb{C}.$$

The cover $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})$ corresponds to choosing a square-root of a .

When $n \geq 3$ the obvious rotationally-symmetric triangulation of (\mathbb{S}, \mathbb{M}) has an associated quiver with potential which consists of a cycle of n arrows, equipped with a nonzero superpotential of degree n . This is known [10, Example 6.7] to be mutation-equivalent to any orientation of the Dynkin diagram of D_n type, necessarily with zero potential.

Example 12.3. *In the non-amenable case $n = 1$ the space $\text{Quad}(\mathbb{S}, \mathbb{M})$ consists of differentials of the form*

$$\phi(z) = (z + c) \frac{dz^{\otimes 2}}{z^2}, \quad c \in \mathbb{C}.$$

The residue at 0 is $\text{Res}_0(\phi) = 4\pi i \sqrt{c}$. The cover $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})$ corresponds to choosing a square-root $s = \sqrt{c}$. Thus

$$\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) \cong \mathbb{C}/\mathbb{Z}_2$$

with \mathbb{Z}_2 acting on \mathbb{C} by $s \mapsto -s$.

There is only one triangulation, and the resulting quiver is a single vertex with no arrows. The category \mathcal{D} is generated by a single spherical object. The mapping class group is trivial, and $\text{MCG}^\pm(\mathbb{S}, \mathbb{M}) = \mathbb{Z}_2$. The group $\text{Aut}_\Delta(\mathcal{D})$ is generated by the shift [1], and the subgroup $\text{Sph}(\mathcal{D})$ by the second shift [2]. The quotient $\text{Stab}_\Delta(\mathcal{D})/\text{Sph}(\mathcal{D})$ is isomorphic to \mathbb{C} with co-ordinate $Z(S)$. Thus the relation

$$\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) \cong \text{Stab}_\Delta(\mathcal{D})/\mathcal{A}ut_\Delta(\mathcal{D}),$$

also holds in this case.

Example 12.4. *In the non-amenable case $n = 2$ the space $\text{Quad}(\mathbb{S}, \mathbb{M})$ consists of differentials of the form*

$$\phi(z) = (az^2 + bz + c) \frac{dz^{\otimes 2}}{z^2}, \quad a \in \mathbb{C}^*, \quad b, c \in \mathbb{C}, \quad b^2 - 4ac \neq 0,$$

modulo the rescaling action of \mathbb{C}^ . Using this we can take $a = 1$; there is then a residual action of \mathbb{Z}_2 acting by $z \mapsto -z$. These differentials have a double pole at $z = 0$ and a fourth order pole at infinity. The respective residues are*

$$\text{Res}_0(\phi) = 4\pi i \sqrt{c}, \quad \text{Res}_\infty(\phi) = 2\pi i b.$$

The cover $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})$ corresponds to choosing a square-root of c . Writing the differential as

$$\phi(z) = (z^2 + 2sz + t^2) \frac{dz^{\otimes 2}}{z^2}, \quad s, t \in \mathbb{C}, \quad s^2 \neq t^2,$$

we see that $\text{Quad}^\pm(\mathbb{S}, \mathbb{M})$ is the quotient $(\mathbb{C}^2 \setminus \Delta)/\mathbb{Z}_2$ where \mathbb{Z}_2 acts by changing the sign of the first co-ordinate s , and $\Delta \subset \mathbb{C}^2$ is the union of the hyperplanes $s = \pm t$. We also have

$$(12.1) \quad \text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) = (\mathbb{C}^2 \setminus \Delta)/\mathbb{Z}_2^{\oplus 2}$$

with the two \mathbb{Z}_2 factors changing the signs of s and t respectively. There is a short exact sequence

$$1 \longrightarrow \mathbb{Z}^{\oplus 2} \longrightarrow \pi_1(\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})) \longrightarrow \mathbb{Z}_2^{\oplus 2} \longrightarrow 1.$$

The mapping class group and its signed version are

$$\text{MCG}(\mathbb{S}, \mathbb{M}) = \mathbb{Z}_2, \quad \text{MCG}^\pm(\mathbb{S}, \mathbb{M}) = \mathbb{Z}_2^{\oplus 2}.$$

There are four tagged triangulations; the two possible taggings of the left-hand triangulation in Figure 36, and the tagged triangulations corresponding to the right-hand picture and its rotation. The corresponding quivers all consist of two vertices with no arrows.




FIGURE 36. Triangulations of a once-punctured disc with two marked points on the boundary.

The category \mathcal{D} is generated by two spherical objects S_1, S_2 lying in the heart of the t -structure corresponding to the triangulation on the left in Figure 36. They have zero Ext groups between them. The twist functor Tw_{S_i} acts by sending S_i to $S_i[2]$ and leaving the other S_j unchanged.

The group of allowable autoequivalences $\mathcal{A}ut_\Delta^{\text{allow}}(\mathcal{D})$ is generated by the spherical twists Tw_{S_i} , together with the autoequivalence swapping S_1 and S_2 , and the shift functor $[1]$. It fits into a short exact sequence

$$1 \longrightarrow \mathbb{Z}^{\oplus 2} \longrightarrow \mathcal{A}ut_\Delta^{\text{allow}}(\mathcal{D}) \longrightarrow \mathbb{Z}_2^{\oplus 2} \longrightarrow 1.$$

It has index 2 in the full group $\mathcal{A}ut_\Delta(\mathcal{D})$, which also contains the element sending S_1 to $S_1[1]$ and leaving S_2 fixed.

The quotient $\text{Stab}_\Delta(\mathcal{D})/\text{Sph}(\mathcal{D})$ is isomorphic to $(\mathbb{C}^*)^2$ with co-ordinates $Z(S_i)$. The isomorphism of Theorem 1.3 is given by

$$Z(S_1) - Z(S_2) = 2\pi i s, \quad Z(S_1) + Z(S_2) = 2\pi i t.$$

In the $Z(S_i)$ co-ordinates the discriminant Δ of (12.1) is given by $Z(S_1)Z(S_2) = 0$. One of the \mathbb{Z}_2 factors acts by exchanging the $Z(S_i)$, and the other acts by changing the signs of the $Z(S_i)$.

12.3. Unpunctured annuli: affine A_n type. Fix integers $p, q \geq 1$ and let (\mathbb{S}, \mathbb{M}) be an annulus whose boundary components contain p and q marked points respectively. This corresponds to differentials on \mathbb{P}^1 with polar type $(p+2, q+2)$. Let $n = p+q$.

The space $\text{Quad}^\pm(\mathbb{S}, \mathbb{M}) = \text{Quad}(\mathbb{S}, \mathbb{M})$ consists of differentials of the form

$$\phi(z) = \prod_{i=1}^n (z - a_i) \frac{dz^{\otimes 2}}{z^{p+2}}, \quad a_i \in \mathbb{C}^*, \quad a_i \neq a_j.$$

These are considered modulo the action of \mathbb{C}^* rescaling z , so if $p \neq q$ we have

$$\text{Quad}(\mathbb{S}, \mathbb{M}) \cong \text{Conf}^n(\mathbb{C}^*)/\mathbb{Z}_q,$$

where $\text{Conf}^n(\mathbb{C}^*)$ denotes the configuration space of n distinct points in \mathbb{C}^* , and the group \mathbb{Z}_q acts by multiplication by a q th root of unity. In the case $p = q$ one should quotient by an extra factor of \mathbb{Z}_2 acting by $z \leftrightarrow 1/z$.

For any triangulation of (\mathbb{S}, \mathbb{M}) the resulting quiver Q is a cycle of n arrows but with a non-cyclic orientation. Thus Q is a non-cyclic orientation of the affine A_{n-1} Dynkin diagram, necessarily with zero potential.

Example 12.5. Consider the non-amenable case $p = q = 1$. The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ parameterizes differentials of the form

$$\phi(z) = (tz + 2s + tz^{-1}) \frac{dz^{\otimes 2}}{z^2}, \quad s \in \mathbb{C}, \quad t \in \mathbb{C}^*, \quad t^2 \neq s^2$$

modulo an action of $\mathbb{Z}_2^{\oplus 2}$, with one generator acting by changing the sign of t , and the other acting trivially, via the automorphism $z \mapsto 1/z$ of \mathbb{P}^1 . Writing $a = s \in \mathbb{C}$ and $b = t^2 \in \mathbb{C}^*$ we obtain

$$\text{Quad}(\mathbb{S}, \mathbb{M}) \cong (\mathbb{C} \times \mathbb{C}^* \setminus \Delta)/\mathbb{Z}_2,$$

where Δ is the hypersurface $b = a^2$ and \mathbb{Z}_2 acts trivially. Write $\text{Quad}'(\mathbb{S}, \mathbb{M})$ for the rigidified moduli space obtained by forgetting about the trivial \mathbb{Z}_2 action. There is a short exact sequence

$$1 \longrightarrow \mathbb{Z} * \mathbb{Z} \longrightarrow \pi_1 \text{Quad}'(\mathbb{S}, \mathbb{M}) \longrightarrow \mathbb{Z} \longrightarrow 1$$

obtained from the obvious projection to \mathbb{C}^* whose fibre over $b \in \mathbb{C}^*$ is $\mathbb{C} \setminus \{\pm\sqrt{b}\}$.

The mapping class group is $\text{MCG}(\mathbb{S}, \mathbb{M}) = \mathbb{Z}_2 \ltimes \mathbb{Z}$ with the \mathbb{Z}_2 acting by exchanging the two boundary components, and \mathbb{Z} acting by a Dehn twist around an equatorial curve. There is a single triangulation of (\mathbb{S}, \mathbb{M}) up to diffeomorphism, and the associated quiver is the Kronecker quiver.




FIGURE 37. Triangulation and quiver for the annulus with one marked point on each boundary component.

The spherical autoequivalence group $\text{Sph}(\mathcal{D})$ is free on the two generators Tw_{S_i} . The mapping class group does not act freely on ordered triangulations, and Proposition 9.9 is false in this case. But the same argument gives a short exact sequence

$$1 \longrightarrow \mathbb{Z} * \mathbb{Z} \longrightarrow \mathcal{A}ut_{\Delta}(\mathcal{D}) \longrightarrow \mathbb{Z} \longrightarrow 1.$$

Theorem 11.2 continues to hold if we replace $\text{Quad}(\mathbb{S}, \mathbb{M})$ by $\text{Quad}'(\mathbb{S}, \mathbb{M})$. The central charges of the two simple objects are given by the elliptic integrals

$$Z(S_i) = \pm 2 \int \sqrt{(tz + 2s + tz^{-1})} \cdot \frac{dz}{z}$$

where the paths of integration are half-loops connecting the two zeroes of ϕ .

12.4. Three-punctured sphere. Let (\mathbb{S}, \mathbb{M}) be the three-punctured sphere. This surface is not amenable, but a version of our Theorem continues to hold. An interesting point is that it seems to be most natural to work with uncompleted Ginzburg algebras in this case (see Remark 9.7).

The space $\text{Quad}(\mathbb{S}, \mathbb{M})$ consists of differentials of the form

$$\phi(z) = \frac{(az^2 + bz + c) dz^{\otimes 2}}{z^2(z-1)^2}$$

for $a, b, c \in \mathbb{C}$ with $b^2 \neq 4ac$, modulo the action of the symmetric group S_3 acting via automorphisms of \mathbb{P}^1 permuting $0, 1, \infty$. The residues at the points $0, 1, \infty$ are

$$\frac{1}{4\pi i} \text{Res}_0 = u = \sqrt{c}, \quad \frac{1}{4\pi i} \text{Res}_1 = v = \sqrt{a+b+c}, \quad \frac{1}{4\pi i} \text{Res}_{\infty} = w = \sqrt{a}.$$

The space $\text{Quad}^{\pm}(\mathbb{S}, \mathbb{M})$ is therefore \mathbb{C}^3 with co-ordinates (u, v, w) , minus the inverse image of the discriminant locus $b^2 - 4ac$, all modulo S_3 acting by permutations on (u, v, w) . The inverse image of the discriminant locus is easily seen to be the divisor $z_1 z_2 z_3 z_4 = 0$, where

$$z_1 = -u + v + w, \quad z_2 = u - v + w, \quad z_3 = u + v - w, \quad z_4 = u + v + w.$$

Thus we conclude that

$$\text{Quad}^\pm(\mathbb{S}, \mathbb{M}) \cong ((\mathbb{C}^*)^3 \setminus \Delta) / S_3$$

where the discriminant Δ is given by $z_1 + z_2 + z_3 = 0$ and the symmetric group Sym_3 acts by permuting (z_1, z_2, z_3) . Hence

$$\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M}) \cong ((\mathbb{C}^*)^3 \setminus \Delta) / (\text{Sym}_3 \ltimes (\mathbb{Z}_2)^{\oplus 3}),$$

where the \mathbb{Z}_2 factors change the signs of u, v, w respectively. The fundamental group then sits in a sequence

$$1 \longrightarrow \pi_1((\mathbb{C}^*)^3 \setminus \Delta) \longrightarrow \pi_1(\text{Quad}_\heartsuit(\mathbb{S}, \mathbb{M})) \longrightarrow \text{Sym}_3 \ltimes \mathbb{Z}_2^{\oplus 3} \longrightarrow 1.$$

The space $(\mathbb{C}^*)^3 \setminus \Delta$ is a trivial \mathbb{C}^* -bundle over its projectivisation, which is the complement of four hyperplanes in \mathbb{CP}^2 . A theorem due to Zariski [38, Lemma, p. 317] asserts that $\pi_1((\mathbb{C}^*)^3 \setminus \Delta)$ is abelian, hence isomorphic to $H_1((\mathbb{C}^*)^3 \setminus \Delta) \cong \mathbb{Z}^{\oplus 4}$.

The mapping class group is $\text{MCG}(\mathbb{S}, \mathbb{M}) = S_3$ and permutes the punctures in the obvious way. The signed mapping class group is $\text{MCG}^\pm(\mathbb{S}, \mathbb{M}) = S_3 \ltimes \mathbb{Z}_2^{\oplus 3}$. There are two triangulations up to the mapping class group action: one has two triangles meeting along three common edges, and the other consists of two self-folded triangles glued along their encircling edges. We claim that the relevant quiver with potential in both cases is the one depicted in Figure 38.




FIGURE 38. The quiver with potential for the 3-punctured sphere.

Note that the associated reduced quiver with potential has three vertices and no arrows. Let \mathcal{A}' be the category of finite-dimensional representations of the ordinary (incomplete) Jacobi algebra of the above quiver with potential, and \mathcal{A} the category of finite-dimensional representations of the completed version. Then $\mathcal{A} \subset \mathcal{A}'$ is the full subcategory of nilpotent representations.

Lemma 12.6. *The category \mathcal{A}' has exactly four indecomposable objects, all of them simple, namely the three vertex simple objects S_1, S_2, S_3 together with the representation S_4 of dimension vector $(1, 1, 1)$ obtained by taking all arrows to be the identity. The subcategory \mathcal{A} is the subcategory consisting of direct sums of the objects S_1, S_2, S_3 .*

Proof. Differentiating the potential we see that the Jacobi algebra has relations

$$a' = bc, \quad b' = ca, \quad c' = ab,$$

which allow us to eliminate a', b', c' . The remaining relations are then

$$(12.2) \quad a = abca, \quad b = bcab, \quad c = cabc.$$

Given a representation of the Jacobi algebra we can split the vector space associated to vertex 1 as $\text{Im}(abc) \oplus \text{Ker}(abc)$. Indeed, by the relations (12.2), if $v = (w)abc$ lies in the kernel of the map abc then $(w)a = (w)abca = (v)a = (v)abca = 0$ and hence $v = 0$. Similar splittings exist at the other vertices and it follows easily from the relations (12.2) that all arrows preserve these splittings. Hence, if E is an indecomposable object of \mathcal{A}' we either have $abc = bca = cba = 0$ or each of these maps is injective.

In the first case it follows from the relations (12.2) that all arrows are 0, and hence, since E is indecomposable it must be one of the vertex simples. In the second case, each of a, b, c is injective, hence they are all isomorphisms, and (12.2) implies that abc is the identity, and similarly for bca and cab . We can then choose the gauge so that a, b, c are identity maps, and since E is indecomposable it follows that the dimension vector must be $(1, 1, 1)$. This unique extra indecomposable S_4 is simple, since there are no maps between it and the vertex simples. \square

It follows from the Lemma that the objects S_i are all spherical and have no extensions between them, since otherwise there would be more indecomposable objects in the category. Consider now the derived category of the uncompleted Ginzburg algebra of the above quiver with potential, and the subcategory \mathcal{D}' of objects with finite-dimensional cohomology. It is a CY_3 triangulated category with a heart $\mathcal{A}' \subset \mathcal{D}$ containing 4 simple, spherical objects, with zero Ext-groups between them.

The Grothendieck group is $K(\mathcal{D}') \cong \mathbb{Z}^{\oplus 4}$ and the dimension vector defines a group homomorphism

$$d: K(\mathcal{D}') \rightarrow \mathbb{Z}^{\oplus 3}.$$

Consider the space $\text{Stab}'(\mathcal{D}') \subset \text{Stab}(\mathcal{D}')$ consisting of those stability conditions whose central charge factors via d . This is acted on by the group of autoequivalences $\text{Aut}'(\mathcal{D}')$ whose action on $K(\mathcal{D}')$ preserves the kernel of d .

The group of all exact autoequivalences of \mathcal{D}' is $\text{Aut}(\mathcal{D}') \cong \text{Sym}_4 \ltimes \mathbb{Z}^{\oplus 4}$, with the four generators ρ_i shifting the four simples S_i respectively, and the symmetric group Sym_4 permuting them. The subgroup $\text{Aut}'(\mathcal{D}')$ contains a subgroup $\text{Sph}(\mathcal{D}') \cong \mathbb{Z}^{\oplus 4}$ generated by the 4 elements ρ_i^2 .

Lemma 12.7. *There is a short exact sequence*

$$1 \longrightarrow \text{Sph}(\mathcal{D}') \longrightarrow \text{Aut}'(\mathcal{D}') \longrightarrow \text{Sym}_3 \ltimes \mathbb{Z}_2^{\oplus 3} \longrightarrow 1.$$

Proof. An element of $\text{Aut}'(\mathcal{D}')$ is determined by its action on the objects S_i , and, up to the action of $\text{Sph}(\mathcal{D}')$, each of these is taken to an object of the form S_j or $S_j[1]$. Consider the transformation

$$\tau_{(12)(34)} : (S_1, S_2, S_3, S_4) \mapsto (S_2, S_1, S_4[1], S_3[1])$$

along with its two conjugates by permutations of (S_1, S_2, S_3) . There is a relation

$$\tau_{(12)(34)} \circ \tau_{(13)(24)} \circ \tau_{(23)(14)} = [1] \in \text{Aut}'(\mathcal{D}')/\text{Sph}(\mathcal{D}').$$

Consider an element σ of the quotient group $\text{Aut}'(\mathcal{D}')/\text{Sph}(\mathcal{D}')$. Composing with the above transformations we can assume that σ takes S_4 to itself. The relation $[S_4] = [S_1] + [S_2] + [S_3]$ then forces σ to be a permutation of the objects (S_1, S_2, S_3) . \square

The space $\text{Stab}'(\mathcal{D}')/\text{Sph}(\mathcal{D}')$ is equal to $(\mathbb{C}^*)^3 \setminus \Delta$ with co-ordinates $z_i = Z(S_i)$, where the discriminant locus Δ is given by $z_1 + z_2 + z_3 = 0$ as before. Thus we obtain an isomorphism

$$\text{Quad}_{\heartsuit}(\mathbb{S}, \mathbb{M}) \cong \text{Stab}'(\mathcal{D}')/\text{Aut}'(\mathcal{D}')$$

which can be thought of as a modified version of Theorem 1.2.

REFERENCES

1. A. Bayer and E. Macri. *The space of stability conditions on the local projective plane*. Duke Math. J. 160: 263–322, (2011).
2. T. Bridgeland. *Stability conditions on triangulated categories*. Ann. of Math. 166:317–345, (2007).
3. T. Bridgeland. *Spaces of stability conditions*. Algebraic Geometry – Seattle 2005. Part I, 1–21. Proc. Sympos. Pure Math. 80, Part 1, Amer. Math. Soc., (2009).
4. T. Bridgeland. *Stability conditions on K3 surfaces*. Duke Math. J. 141 (2008), no. 2, 241291. Duke Math. J. (2009).
5. T. Bridgeland. *Stability conditions on a non-compact Calabi-Yau threefold*. Comm. Math. Phys. 266:715–33, (2006).
6. G. Cerulli Irelli, B. Keller, D. Labardini-Fragoso and P.-G. Plamondon. *Linear independence of cluster monomials for skew-symmetric cluster algebras*. Preprint, available at arXiv:1203.1307.
7. G. Cerulli Irelli and D. Labardini-Fragoso. *Quivers with potential associated to triangulated surfaces, Part III: Tagged triangulations and cluster monomials*. Preprint, to appear in Compositio Math., available at arXiv:1108.1774.
8. H. Derksen, J. Weyman and A. Zelevinsky. *Quivers with potential and their representations, I: Mutations* Selecta Math. (N.S.) 14:59–119, (2008).
9. A. Fletcher and V. Markovich, *Quasiconformal maps and Teichmüller theory*. Oxford Graduate Texts in Mathematics, 11. Oxford University Press, Oxford, 2007. viii+189 pp.
10. S. Fomin, M. Shapiro and D. Thurston. *Cluster algebras and triangulated surfaces, Part I: cluster complexes*. Acta Math. 201:83–146, (2008).
11. S. Fomin and D. Thurston. *Cluster algebras and triangulated surfaces part II: Lambda lengths*. Preprint, available at arXiv:1210.5569.
12. C. Geiss, D. Labardini-Fragoso and J. Schröer. *Preprint in preparation*.
13. D. Gaiotto, G. Moore and A. Neitzke. *Four-dimensional wall-crossing via three-dimensional field theory*. Comm. Math. Phys. 299:163–224, (2010).
14. D. Gaiotto, G. Moore and A. Neitzke. *Wall-crossing, Hitchin systems and the WKB approximation*. Preprint, available at arXiv:0907.3987.

15. V. Ginzburg. *Calabi-Yau algebras*. Preprint, available at arXiv:math/0612139.
16. Weiwen Gu. *Graphs with non-unique decomposition and their associated surfaces*. Preprint, available at arXiv:1112.1008.
17. D. Happel, I. Reiten and S. Smalø. *Tilting in abelian categories and quasitilted algebras*. Mem. Amer. Math. Soc. 120, no. 575, (1996).
18. B. Keller. *On differential graded categories*. International Congress of Mathematicians. Vol. II, 151190, Eur. Math. Soc., Zürich, (2006).
19. B. Keller. *On cluster theory and quantum dilogarithm identities*. Preprint, available at arXiv:1102.4148.
20. B. Keller and D. Yang. *Derived equivalences from mutations of quivers with potential*. Adv. Math. 226:2118–2168, (2011).
21. A. D. King. *Moduli of representations of finite-dimensional algebras*. Quart. J. Math. 45:515–530, (1994).
22. A. D. King and Yu Qiu. *Exchange graphs of acyclic Calabi-Yau categories*. Preprint, available at arxiv 1109.2924v2.
23. M. Kontsevich and Y. Soibelman. *Stability structures, motivic Donaldson-Thomas invariants and cluster transformations*. Preprint, available at arXiv:0811.2435.
24. D. Labardini-Fragoso. *Quivers with potentials associated to triangulated surfaces*. Proc. Lond. Math. Soc. 98:797–839, (2009).
25. D. Labardini-Fragoso. *Quivers with potentials associated to triangulated surfaces, Part II: Arc Representations*. Preprint, available at arXiv:0909.4100.
26. D. Labardini-Fragoso. *Quivers with potentials associated to triangulated surfaces, part IV: Removing boundary assumptions*. Preprint, available at arXiv:1206.1798.
27. H. Masur. *Closed trajectories for quadratic differentials with an application to billiards*. Duke Math. J. 53:307–314, (1986).
28. H. Masur and A. Zorich. *Multiple saddle trajectories on flat surfaces and the principal boundary of the moduli spaces of quadratic differentials*. Geom. Funct. Anal. 18:919–987, (2008).
29. A. Papadopoulos. *Metric Spaces, Convexity and Nonpositive Curvature*. European Mathematical Society, (2005).
30. P. Seidel, R.P. Thomas. *Braid group actions on derived categories of coherent sheaves*. Duke Math. J. 108:37–108, (2001).
31. S. Smale. *Diffeomorphisms of the 2-sphere*. Proc. Amer. Math. Soc. 10:621–626, (1959).
32. I. Smith. *Quiver algebras as Fukaya categories*. Preprint in preparation, 2013.
33. K. Strebel. *Quadratic differentials*. Ergebnisse der Mathematik und ihrer Grenzgebiete, Springer-Verlag, Berlin, 1984.
34. T. Sutherland. *The modular curve as the space of stability conditions of a CY3 algebra*. Preprint, available at arXiv:1111.4184.
35. T. Sutherland. *Stability conditions and Seiberg-Witten curves*, Ph.D. Thesis, University of Sheffield. To appear, 2013.
36. W. Veech. *Moduli spaces of quadratic differentials*. J. Analyse Math. 55:117–171, (1990).
37. J. Woolf. *Stability conditions, torsion theories and tilting*. J. Lond. Math. Soc. 82:663–682, (2010).
38. O. Zariski. *On the Problem of Existence of Algebraic Functions of Two Variables Possessing a Given Branch Curve*. Amer. J. Math. 51: 305328, (1929).

TOM BRIDGELAND, ALL SOULS COLLEGE, OXFORD, OX1 4AL, ENGLAND.

IVAN SMITH, CENTRE FOR MATHEMATICAL SCIENCES, CAMBRIDGE, CB3 0WB, ENGLAND.